

Article

Automatic Detection and Distinction of Retinal Vessel Bifurcations and Crossings in Colour Fundus Photography

Harry Pratt ^{1,*}, Bryan M. Williams ¹, Jae Yee Ku ^{1,2}, Charles Vas ¹, Emma McCann ¹, Baidaa Al-Bander ³, Yitian Zhao ⁴, Frans Coenen ² and Yalin Zheng ^{1,5}

¹ Department of Eye and Vision Science, Institute of Ageing and Chronic Disease, University of Liverpool, Liverpool L69 3BX, UK; bryan@liverpool.ac.uk (B.M.W.); Jae.Ku@liverpool.ac.uk (J.Y.K.); hlcvas@liverpool.ac.uk (C.V.); hlemccan@liverpool.ac.uk (E.M.); yzheng@liverpool.ac.uk (Y.Z.)

² Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK; coenen@liverpool.co.uk

³ Department of Electrical Engineering, University of Liverpool, Liverpool L69 3BX, UK; hsbalban@liverpool.ac.uk

⁴ Cixi Institute of Biomedical Engineering, Ningbo Institute of Industrial Technology, Chinese Academy of Sciences, Ningbo 315201, China; Yitian.Zhao@liverpool.ac.uk

⁵ St Paul's Eye Unit, Liverpool Royal University Hospital, Liverpool L69 3BX, UK

* Correspondence: h.pratt@liverpool.ac.uk

† Current address: Department of Eye and Vision Science, Institute of Ageing and Chronic Disease, University of Liverpool, Liverpool L7 8TX, UK.

Received: 7 November 2017; Accepted: 14 December 2017; Published: 22 December 2017

Abstract: The analysis of retinal blood vessels present in fundus images, and the addressing of problems such as blood clot location, is important to undertake accurate and appropriate treatment of the vessels. Such tasks are hampered by the challenge of accurately tracing back problems along vessels to their source. This is due to the unresolved issue of distinguishing automatically between vessel bifurcations and vessel crossings in colour fundus photographs. In this paper, we present a new technique for addressing this problem using a convolutional neural network approach to firstly locate vessel bifurcations and crossings and then to classifying them as either bifurcations or crossings. Our method achieves high accuracies for junction detection and classification on the DRIVE dataset and we show further validation on an unseen dataset from which no data has been used for training. Combined with work in automated segmentation, this method has the potential to facilitate: reconstruction of vessel topography, classification of veins and arteries and automated localisation of blood clots and other disease symptoms leading to improved management of eye disease.

Keywords: medical image analysis; machine learning; convolutional neural networks; retinal imaging; retinal vessels; fundus photography; vessel classification

1. Introduction

Vascular conditions present a challenging public health problem as they become more common due to global ageing [1]. Vascular conditions are often life-threatening and blood vessel damage caused from common health issues such as diabetes, hypertension and strokes can lead to significant health complications. It is, therefore, of great importance to better understand and be able to manage such conditions. The retina is the only inner organ which can be directly imaged, using a fundus camera, and also serve as a “window” for the diagnosis of systematic diseases such as: cerebral malaria, stroke, dementia and cardiovascular diseases [2]. It is also significant that pathologies often affect veins and arteries differently. For example, in diabetic retinopathy, abnormalities typically occur in veins such as venous beading which is a significant predictor to the sight-damaging proliferative stage of the

condition. With the availability of imaging techniques such as colour fundus photography, fundus angiography and recent optical coherence tomography angiography, there is a significant need for automated vessel analysis techniques [3,4].

There has been a considerable amount of work, in recent years, aimed at the effective segmentation of retinal blood vessels in fundus photography, which is a prerequisite step for blood vessel analysis. Work such as [3–5] has been able to achieve increasingly improved segmentation of retinal vessels. However, a significant remaining challenge is to distinguish between vessel bifurcations and vessel crossings. A vessel bifurcation is where a mother vessel branches into two daughter vessels, while a vessel crossing is where one vessel passes over another but does not connect to it. This is important for tracking vessels, separating veins from arteries and providing for quantitative analysis of vasculature.

For example, we must be able to trace back along the vessel when a blood clot has been identified. The current inability to accurately identify vessel crossings after or during vessel segmentations hinders this. It is also important to monitor progress of a vessel after vein and artery occlusions; being able to identify and distinguish vessel crossings and bifurcations facilitates this. Automating the detection and classification of vessel bifurcations and crossings also allows us to aid clinicians in detecting vascular abnormalities. The vast amount of vessels and vessel junctions within the retina make this task a laborious one for clinicians; by automating the process we can save time for treatment while maintaining accuracy. The vasculature can be obtained through vessel tracking or pixel-based classification. Detecting bifurcations and crossings are critical to either of these vasculature reconstruction methods. The detected and classified vessel junctions can be used in combination with vessel segmentation, or used in vessel tracking methods to detect the source of irregular vasculature.

The previous work on vessel bifurcations and junctions has involved using orientation scores to detect bifurcations and junctions in retinal images [6]. In contrast to the method we proposed, which is a fully automated system that uses only the image to determine the diagnosis, the work in [6] required 24 orientation processes for each image before training. However, the results in the paper show that the features within the image are extractable. There has also been similar orientation-based work in [7,8].

For applications in image analysis and classification, Convolutional Neural Networks (CNNs), a branch of deep learning, has achieved state of the art results for many problems. The 1970s saw the introduction of network architectures being used to analyse image data [9]. These had useful applications and allowed challenging tasks, such as handwritten character recognition [10], to be achieved. Decades later, there were several breakthroughs in neural networks that lead to vast improvements in their implementation, such as the introduction of dropout [11] and rectified linear units [12]. These theoretical enhancements and the accompanying increase in computing power through graphical processor units (GPUs) meant that CNNs became viable for more complex image recognition problems. Presently, large CNNs are used to successfully tackle highly complex image recognition tasks with many object classes to an impressive standard [13,14]. The recent improvements in image recognition problems mentioned present an opportunity for more efficient and accurate methods of our vessel problem. CNNs are used in many of the current state-of-the-art image classification tasks including medical imaging. Hence, we use this method combined with expert segmented fundus images and skeletonisation [15,16] to detect and classify vessel bifurcations and crossings within fundus images.

There are many different architectures for neural networks. Recently residual networks have achieved impressive results on the highly competitive competition of ImageNet detection, ImageNet localisation, COCO detection, and COCO segmentation [17]. They were then widely used in the following 2016 ImageNet competition due to their impressive performance on general large data sets of small images; such as the MNIST [18] dataset for handwritten digits 0–9 and CIFAR-10 [14], a dataset of 10 classes of colour images. This network learns from the residual of the identity of the previous layer of a new layer in order to learn features more effectively. This makes the network ideal for our patch-based method as the higher level features can distinguish between the background of the retina

and the bifurcations and crossings. Hence, the Res18 network structure, containing 18 residual layers, was used in the CNNs throughout this paper.

In this paper, we present a new hierarchical approach, that utilises deep learning, to first automatically determine the locations of blood vessel bifurcations and crossings in colour fundus images, and then to distinguish between vessel bifurcation and crossings. We employ an available segmentation of the vessel structure, although an automatic segmentation procedure could be incorporated, to identify points along blood vessels. Annotated image datasets with identified vessel bifurcations and crossings aid our deep learning framework as we use a supervised learning method to solve this image recognition problem. From the original fundus images of the DRIVE dataset we created small patches of images using a skeletonisation of the vessels. We use a convolutional neural network approach which is trained on some of the patches of the fundus images using the expert ground truth for optimization. A matching network architecture is then used and trained to learn new convolution filters to distinguish between vessel bifurcations and crossings. The results is a novel method which is capable of identifying and classifying vessel bifurcations and crossings without user intervention.

The rest of this paper is organised as follows. In Section 2, we present our new automatic approach for locating and identifying crossings and bifurcations of retinal vessels, in Section 3 we demonstrate that proposed method yields robust, state-of-the-art results and in Sections 4 and 5 we present our conclusions and discuss future work. This paper is an extension of the paper [19] extended and more substantial results and a refined method for accuracy. The figures used are cited throughout.

2. Methods

Firstly we identifying patches of fundus images $z(x)$. During our experiments we found that the optimal size for both performance and collection of the patches was 21 by 21 pixels. All of the patches used throughout this method were of this size. We make use of available vessel segmentations given as binary functions defined on the domain, and perform a skeletonisation process on this domain. The patches are then produced along the skeleton so that each contains some of vessel structure. Furthermore, after creating the patches we train our Res18 convolutional neural network to identify the patches which include either a bifurcation or a crossings. The res18 neural network contains 18 convolutional layers learned using the residual of the previous convolutional layer as in [17]. The network contains 11,181,570 trainable weight parameters for optimization and the architecture layout can be found in the supplementary material. Another network with the same architecture is then trained on the patches that have been graded to have bifurcations and crossings to distinguishing the type of vessel junction located.

We tested the ability of our algorithm using 40 images from the DRIVE database with manual segmentations [4]. We also studied the variability between grading and how this relates to the trained network for each grader. The data split was 30 images for training the neural networks, leaving 10 for testing. While this may seem a small number for a machine learning approach, our patch-based method means that the images generated for training numbered more than 100,000 providing sufficient data. Ground truth annotations of vessel crossings and bifurcations were provided by two graders (G1 and G2).

2.1. Datasets

The images used to implement our framework are from the Digital Retinal Images for Vessel Extraction (DRIVE) database with manual segmentations [4]. The images in the DRIVE dataset were obtained from a diabetic retinopathy screening program in The Netherlands. The images were acquired using a Canon CR5 non-mydratic 3CCD camera with a 45 degree field of view (FOV) using 8 bits per colour plane at 768 by 584 pixels.

Moreover, we use the IOSTAR dataset [6,20] for testing the robustness of the method. Our networks are trained on the DRIVE Dataset, but all of them are tested on the unseen IOSTAR

dataset for further validation. The IOSTAR dataset is made of the images taken with EasyScan camera (provided by i-Optics B.V., The Hague, The Netherlands). The original images have a resolution of 1024 by 1024 (14 $\mu\text{m}/\text{px}$), and a 45 degree field of view. For the moment, vessels, bifurcations and crossings of 24 images have been annotated and corrected by two different experts, the same experts that graded the DRIVE dataset. For testing on the IOSTAR dataset, which has the same field of view as DRIVE, the images were resized using bilinear interpolation to the dimensions of the DRIVE images. Patches were then extracted in the same way with both datasets to allow for fair comparison purposes. This process can be used to compare with any dataset of varying image size.

The datasets were graded separately by 3 expert graders to compare variability between the networks and between the graders themselves. The graders labelled bifurcations by clicking as close to the centre of the junction as possible. This allowed for it to be a simple operation so that the graders focus could remain on image.

2.2. Skeletonisation and Patch Extraction

We consider patches of the fundus images centred along the segmented vessels. In order to restrict the number of patches for training to a manageable number, and reduce bias, we aim to reduce to segmentation of the vessels to a skeleton and consider regions centred only on these points. We achieve this by performing a skeletonisation of the level set function $\phi(\mathbf{x})$ for each image.

$$\kappa_1^j(\alpha^1) = r_j \begin{pmatrix} 0 & 0 & 0 \\ \alpha_1^1 & 1 & \alpha_2^1 \\ 1 & 1 & 1 \end{pmatrix}, \quad \kappa_2^j(\alpha^2) = r_j \begin{pmatrix} \alpha_1^2 & 0 & 0 \\ 1 & 1 & 0 \\ \alpha_2^2 & 1 & \alpha_3^2 \end{pmatrix}, \quad (1)$$

We convolve the level set function with the kernels in Equation (1) where r_j denotes rotation of the matrix by a multiple j of $\pi/2$ radians and $\alpha^1 = (\alpha_1^1, \alpha_2^1)^\top \in \Psi^2$, $\alpha^2 = (\alpha_1^2, \alpha_2^2, \alpha_3^2)^\top \in \Psi^3$ where $\Psi = \mathbb{Z} \cap [0, 1]$. These kernels are shown in Figure 1. We thin the segmentation of the vessels by removing the points which are centred on regions matching the above filters. That is, we set such points as background points.

$$\varphi^{\ell+1} = \mathcal{F}_{i,j}(\varphi^\ell), \quad \ell = 0, 1, \dots, \quad \varphi_0(\mathbf{x}) = \phi(\mathbf{x})$$

$$\mathcal{F}_{i,j}(\varphi) = \varphi - 1 + H \left(\left(\kappa_i^j(\alpha^i) * \varphi - \sum \kappa_i^j(\alpha^i) \right)^2 \right), \quad (2)$$

We achieve this by iterating as in Equation (2), beginning with $l_0^1 = 0$ and cycling through $i \in \{1, 2\}$, $j \in \{0, 1, 2, 3\}$.




Figure 1. Kernel functions for skeletonisation [19] Reproduced with permission.

Following this, we extract the patches by cropping the image $z(\mathbf{x})$ to 21×21 pixel windows $\Theta_{\mathbf{p}}$ centred on points $\mathbf{p}$ in the set Y of points considered the foreground of the skeletonised vessel map.

The patch size was selected so that bifurcations, and crossings and branches, in the vessel would fit within one patch. The patches are given by:

$$\Theta_{\mathbf{p}} = \{\mathbf{q} \in \Omega \mid |\mathbf{p} - \mathbf{q}| \leq 10\}, \quad \mathbf{p} \in Y = \{\mathbf{p} \in \Omega \mid \varphi(\mathbf{p}) = 1\}.$$

In the training stage, the set of patches (Θ) of the images in the training set are used to train the neural network to identify whether a bifurcation or crossing is contained in the image patch. In the test stage, the trained CNN classifies the patches accordingly. This step is described below.

2.3. Junction Distinction—CNN $\mathbb{C}_1$

To identify the vessel bifurcations and crossings within the patches created we train our CNN on a high-end Graphics Processor Unit (GPU). The large random access memory of the Nvidia K40c means that we were able to train on the whole dataset of patches at once. The Nvidia K40c contains 2880 CUDA cores and comes with the Nvidia CUDA Deep Neural Network library (cuDNN) for GPU learning. The deep learning package Keras [21] was used alongside the Theano machine learning back end to implement the network. After training, the feed forward process of the CNN can classify the patches produced from a single image in under a second.




Figure 2. Example outcomes of first part of algorithm: locating bifurcations and crossings. (a) Fundus Image $z(\mathbf{x})$; (b) Vessel Map $\phi(\mathbf{x})$; (c) Skeletonisation $\phi(\mathbf{x})$; (d) Patch Boundaries; (e) Patch Classification; (f) Junction Location. Reproduced with permission from [19].

We used the Res18 network architecture [17] as deep levels of convolution were required to distinguish the vessel junction type in our small patches. The residual layers incorporate activation, batch normalisation, convolutional, dense and maxpooling layers. We also use L^2 regularisation to improve weight training. There were approximately 100,000 patches for training and 30,000 for testing

in the junction distinction problem. The classes were weighted as a ratio of junction to background due to the fact that bifurcations and crossings in the training and testing patches were sparse at a ratio of 1:39. The network was optimised using Adam stochastic optimisation for backpropagation [22]. The network was trained to classify the patches to give a binary classification of either vessel junction or background. Gaussian initialisation was used within the network to reduce initial training time. The loss function used for the optimisation was the widely used categorical cross-entropy function. Training was undertaken until reduction of the loss plateaued to obtain optimal results. These results can be seen in Figure 2.

2.4. Locate the Centres

Following the neural network classification, which tell us if a bifurcation or crossing is contained within a patch, we aim to find the locations of the points.

$$t(\mathbf{q}) = \sum_{\mathbf{p} \in Y} s^{\mathbf{p}}(\mathbf{q}, l), \quad s^{\mathbf{p}}(\mathbf{q}, l) = \begin{cases} l_{\mathbf{p}}^1 & \text{if } q \in \Theta(\mathbf{p}) \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

We achieve this by forming the cumulative sum image shown in Equation (3) and taking the local maxima $\mathbf{r} \in Y$ as points of interest. We then aim to determine whether points are at crossings or bifurcations.

2.5. Junction Classification $\mathbb{C}_2$

We extract the patches $\Theta(\mathbf{r})$ and use these to train a neural network to distinguish between crossings and bifurcations as shown in Figure 3. The second neural network was trained with the Res18 architecture, like the first. Using a relatively small training set of patches, as from our images the majority of patches did not contain bifurcations and crossings, we trained our network in similar fashion to that used in the previous step. Weighted classes were introduced again to cater for the imbalance, in that images from the bifurcation class were substantially more prominent than that of the cross class.

Depending on the patch method there were around 800–2500 patches containing a junction that was used for training which can be seen in Figure 4. In all methods there were approximately twice as many junction patches containing bifurcation vessels compared to patches containing vessels crossing. Training was performed until a plateau in the reduction of the loss function was reached indicating no further improvement.




Figure 3. Example outcomes of second part of algorithm: classifying bifurcations and crossings as bifurcations and crossings. (a) Identified bifurcations and crossings; (b) bifurcation Points; (c) Vessel Crossings. Reproduced with permission from [19].




Figure 4. Example of $\mathbb{C}_2$ input. Rows 1 and 2 (resp. 3 and 4): training patches with crossings (resp. bifurcations) and their enhanced counterparts for presentation. The neural networks were able to achieve good results using the patches without enhancement. Reproduced with permission from [19].

3. Results

We present our results on a patch by patch basis as well as in the fundus image form with vessel bifurcations and splittings labeled. The patch detection and classification information is then used to in a probability map type reconstruction of the fundus image to produce the final appropriate vessel bifurcations and splittings, as demonstrated in Figures 5 and 6. Here we present both the patch accuracy results and the final classified and vessel type distinguished images. We measure sensitivity, specificity and accuracy of the final image-based result as follows. Since the region of a junction is not restricted to a single point, we allow a region $r(x, y)$ of 10 pixels either side of an annotated point at (x, y) to be considered the correct region. That is, we split each image domain Ω into two sets

$$\Omega_1 = \{(x, y) | j \in r(x, y)\}, \quad \Omega_2 = \Omega \setminus \Omega_1$$

where j denotes a junction point location. Ω_1 is considered the true (junction) set and Ω_2 is considered the background set. We then calculate error measures based on this and report the mean measures.

For validation we used the test images from the DRIVE dataset. Furthermore, we used the separate IOSTAR dataset, and another expert grader (G3), for testing and comparison of the patch detection and classification method. We show in Tables 1 and 2 the results of training our neural networks on the data provided by graders 1, 2 and 3. In each case, we use patches extracted from 30 images from the DRIVE dataset to train our network. This relates to 101,416 patches of vessel junctions and 216,756 other patches. From the 101,416 patches there are 67,650 patches of vessel crossings and 33,766 patches of vessel bifurcations. This network is then tested by using our network to classify patches from the remaining ten images of the DRIVE dataset. This relates to 72,980 non-vessel patches and 31,026 vessel patches of which 9176 are vessel crossings and 21,850 are vessel bifurcations. We compare the results with the annotations provided by the grader in question, achieving high accuracies of 0.76, 0.76 and 0.77 for graders 1, 2 and 3 respectively. Furthermore, we use the trained network to classify images from the IOSTAR dataset, comparing with the annotations provided. With each trained network, the accuracy is lower for this unseen dataset, but the sensitivity is retained. The IOSTAR dataset gave us 132,064 non-vessel junction patches and 52,878 vessel junction patches with 14,228 vessel crossings and 38,650 vessel bifurcations. Following the detection, we resolve the

patch-based results into the original images in order to identify individual junctions and to measure the detection performance for each image. Tables 3 and 4 shows the results obtained from the networks trained by the first annotations of grader 2 and grader 3 and tested on the 10 remaining test images of the DRIVE dataset. The results were compared with the annotations of grader 1 (G1A1), the first and second annotations of grader 2 (G2A1 and G2A2) as well as the first and second annotations of grader 3 (G3A1 and G3A2). Excellent performance of 81% is achieved for the network trained and tested on grader 2's first annotation. The result is similar when comparing to the other graders. Good performance of 74% is also achieved for the network trained and tested on grader 3's first annotation with similar, and even improved, results when comparing to other annotations. We retain good accuracy for the classification task, as shown in Table 5 and Figures 7 and 8, achieving accuracies of ≥ 0.70 for distinguishing between detected vessel crossings and bifurcations. The results were a little lower for the IOSTAR dataset but this, and the detection results, may be improved by including some of this data in the training of the networks.

Table 1. CNN-based Detection results.

Training on Grader 1			
Test Set	Accuracy	Sensitivity	Specificity
G1A2	0.7622	0.7398	0.7714
IOSTAR	0.7349	0.5225	0.8199
G1A1	0.8055	0.6244	0.8688
Grader 2			
Test Set	Accuracy	Sensitivity	Specificity
G2A2	0.7620	0.7472	0.7681
IOSTAR	0.6302	0.7639	0.5767
G1A1	0.7537	0.7704	0.7479
Grader 3			
Test Set	Accuracy	Sensitivity	Specificity
G2A2	0.7654	0.7408	0.7756
IOSTAR	0.6586	0.7466	0.6234
G1A1	0.7574	0.7664	0.7543

Table 2. Confusion Matrices for CNN-based Detection results. BG denotes background, JC denotes junctions. True labels are along rows, predicted along columns.

Training on Grader 1								
G1A2	BG	JC	IOSTAR	BG	JC	G1A1	BG	JC
BG	56,732	7926	BG	108,284	25,248	BG	66,969	10113
JC	16,811	22,537	JC	23,780	27,630	JC	10,115	16,809
Grader 2								
G2A2	BG	JC	IOSTAR	BG	JC	G1A1	BG	JC
BG	56,847	7701	BG	76,161	12,484	BG	57,651	6181
JC	17,056	22,762	JC	55,903	40,394	JC	19,433	20,741
Grader 3								
G2A2	BG	JC	IOSTAR	BG	JC	G1A1	BG	JC
BG	57,039	7896	BG	82,329	13,399	BG	58,143	6290
JC	16,504	22,567	JC	49,735	39,479	JC	18,941	20,632

Table 3. Image-based Detection results.

Grader 2					
Tested on	G1A1	G2A1	G2A2	G3A1	G3A2
Sensitivity	0.9300	0.8474	0.8474	0.8923	0.9023
Specificity	0.7836	0.8090	0.8024	0.7991	0.7998
Accuracy	0.7968	0.8148	0.7157	0.8114	0.8129

Grader 3					
Tested on	G1A1	G2A1	G2A2	G3A1	G3A2
Sensitivity	0.9869	0.9354	0.9637	0.9634	0.9731
Specificity	0.7603	0.7227	0.7135	0.7100	0.7107
Accuracy	0.7199	0.7516	0.7436	0.7388	0.7405

Table 4. Confusion Matrices for CNN-based Classification results. BF denotes bifurcations, CR denotes crossings. True labels are along rows, predicted along columns.

Training on Grader 1					
G1A2	BF	CR	IOSTAR	BF	CR
BF	13,289	2068	BF	12,375	1907
CR	2379	4801	CR	7821	5527

Grader 2								
G2A2	BF	CR	IOSTAR	BF	CR	G1A1	BF	CR
BF	12,250	3158	BF	14,430	3262	BF	11,468	2547
CR	3780	3574	CR	15,095	7607	CR	2951	3775

Grader 3								
G3A2	BF	CR	IOSTAR	BF	CR	G1A1	BF	CR
BF	12,046	2585	BF	15,475	2820	BF	11,414	2387
CR	3864	4072	CR	13,381	7803	CR	2929	3902

Table 5. CNN-based Classification results.

Grader 1			
Test Set	Accuracy	Sensitivity	Specificity
G1A2	0.8027	0.6989	0.8482
IOSTAR	0.6479	0.7435	0.6127

Grader 2			
Test Set	Accuracy	Sensitivity	Specificity
G2A2	0.6952	0.5309	0.7642
IOSTAR	0.5456	0.6999	0.4887
G1A1	0.7349	0.5971	0.7953

Grader 3			
Test Set	Accuracy	Sensitivity	Specificity
G3A2	0.7142	0.6117	0.7571
IOSTAR	0.5896	0.7345	0.5363
G1A1	0.7423	0.6204	0.7958




Figure 5. Example of identifying bifurcations and crossings in fundus images. (a), (b), (d) are left eye fundus images and (c), (e) and (f) are right eye fundus images. Reproduced with permission from [19].




Figure 6. Both (a) and (b) examples from the test set that have been run through the algorithm. Here we demonstrate how the patch classification leads to the building up of a vessel map from identifying and classifying the vessel junctions and reconstructing the classified patches. The detected junctions are shown on the fundus image showing that the algorithm clearly identifies junction points. Reproduced with permission from [19].




Figure 7. (a) and (c) show the identified bifurcations (a) and crossings (c) for an example image from the DRIVE dataset. The results are shown along with the annotations provided by each grader. The annotation of grader 1 is shown by a red x, grader 2 by a blue x and green o, grader 3 by a cyan x and black o; (b) and (d) are zoomed in to demonstrate the negligible difference in the classification of the vessel bifurcation (b) and crossing (d) from the grader's annotations and the consistency in the annotations provided. Reproduced with permission from [19].




Figure 8. Cont.




Figure 8. Example of distinguishing between crossings and bifurcations in fundus images. Rows one shows the classified bifurcations in (a–c) and row two the respective crossings in (d–f). Likewise, row three shows the classified bifurcations in (g–i) and row four the respective crossings in (j–l). Reproduced with permission from [19].

4. Discussion

We have produced a method that can learn to detect and classify vessel bifurcations and crossings using a very small dataset of 40 fundus images that had been manually classified for bifurcations and crossings and their type. Using the CNN $\mathbb{C}_1$, we managed to detect the bifurcations and crossings to an impressive detection accuracy of over 90% due in part to the relatively large amount of patches containing bifurcations and crossings. Along with the skeletonisation, our deep learning classification $\mathbb{C}_2$ for vessel type gave us a high accuracy. The classifications statistics are similar to that of the intravariability between graders and hence the method has more chance of reaching higher results should a consistent grading be given within the patches. This can be seen by the higher testing results on the third grader. Increasing the size of our dataset would allow better distinction in the classification of the vessel bifurcations and crossings. It is worth noting that junction type training was undertaken on a couple of thousand patches and tested on around 800. Through training on more images the model could be fine tuned to refine the filters and increase distinction accuracy.

The current algorithm works well for images which have been manually segmented but this time-consuming task could be further extended to incorporate automatic segmentation techniques [3]. A further very useful extension would be to automatically determine whether the artery or vein is in front given arteriovenous crossings, along with consideration of intra and inter-observer variability. In order to better identify and classify bifurcations and crossings with other nearby bifurcations and crossings, it would be useful to consider extending our method to a multi-scale approach. Furthermore, while we have included images with different vessel pathologies due to diseases such as retinopathy, using datasets containing other retinal diseases this could be studied further to investigate how this affects the detection and classification of vessel junctions.

5. Conclusions

The challenging task of detecting and classifying vessel bifurcations and crossings in fundus images is achieved to a high level of accuracy using our method. The variability over datasets without training on the dataset represents a robust algorithm for unseen images. The ability to expand on this method to make the detection both quicker and more accurate than manual classification is possible. These preliminary results demonstrate that the overall framework, including the deep learning approach proposed, is a viable technique to accurately find and identifying vessel bifurcations and crossings with little training data. More extensive testing of this framework could be undertaken to assess the transferability of these results and patch sizes to different size images from different datasets. However, there is no reason why this framework would not be directly applicable to another dataset.

Supplementary Materials: The following are available online at www.mdpi.com/2313-433X/4/1/4/s1. The supplementary material shows the architecture of the neural network used in this work. The arrows indicate the links between layers and the network runs from the input at the top of the image to the output at the bottom of the image.

Acknowledgments: Harry Pratt acknowledges PhD studentship from Fight for Sight, UK. This project is funded in part by the National Institute for Health Research's i4i Programme. This paper summarises independent research funded by the National Institute for Health Research (NIHR) under its i4i Programme (Grant Reference Number II-LA-0813-20005). The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health.

Author Contributions: Harry Pratt and Bryan M. Williams conceived, designed and implemented the experiments. Jae Yee Ku gave clinical input on the problem. Charles Vas and Emma McCann performed the grading of the data. Baidaa Al-Bander and Yitian Zhao analyzed the data. Frans Coenen and Yalin Zheng contributed to analysis and paper writing.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Benjamin, E.J.; Blaha, M.J.; Chiuve, S.E.; Cushman, M.; Das, S.R.; Deo, R.; de Ferranti, S.D.; Floyd, J.; Fornage, M.; Gillespie, C.; et al. Heart Disease and Stroke Statistics 2017 Update: A Report From the American Heart Association. *Circulation* **2017**, *135*, e146–e603.
2. MacGillivray, T.; Trucco, E.; Cameron, J.; Dhillon, B.; Houston, J.; Van Beek, E. Retinal imaging as a source of biomarkers for diagnosis, characterization and prognosis of chronic illness or long-term conditions. *Br. J. Radiol.* **2014**, *87*, 20130832.
3. Zhao, Y.; Rada, L.; Chen, K.; Harding, S.P.; Zheng, Y. Automated vessel segmentation using infinite perimeter active contour model with hybrid region information with application to retinal images. *IEEE Trans. Med. Imaging* **2015**, *34*, 1797–1807.
4. Staal, J.; Abramoff, M.D.; Niemeijer, M.; Viergever, M.A.; Van Ginneken, B. Ridge-based vessel segmentation in color images of the retina. *IEEE Trans. Med. Imaging* **2004**, *23*, 501–509.
5. Chutatape, O.; Zheng, L.; Krishnan, S.M. Retinal blood vessel detection and tracking by matched Gaussian and Kalman filters. In Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Hong Kong, China, 1 November 1998; Volume 6, pp. 3144–3149.

6. Abbasi-Sureshjani, S.; Smit-Ockeloen, I.; Bekkers, E.; Dashtbozorg, B.; ter Haar Romeny, B.M. Automatic detection of vascular bifurcations and crossings in retinal images using orientation scores. In Proceedings of the 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI), Prague, Czech Republic, 13–16 April 2016.
7. Dashtbozorg, B.; Abbasi-Sureshjani, S.; Zhang, J.; Huang, F.; Bekkers, E.J.; ter Haar Romeny, B.M. Infrastructure for Retinal Image Analysis. In Proceedings of the Ophthalmic Medical Image Analysis Third International Workshop (OMIA 2016), Athens, Greece, 21 October 2016.
8. Bekkers, E.; Duits, R.; Berendschot, T.; ter Haar Romeny, B. A multi-orientation analysis approach to retinal vessel tracking. *J. Math. Imaging Vis.* **2014**, *49*, 583–610.
9. Fukushima, K. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* **1980**, *36*, 193–202.
10. Cun, Y.L.; Boser, B.; Denker, J.S.; Howard, R.E.; Habbard, W.; Jackel, L.D.; Henderson, D. *Advances in Neural Information Processing Systems 2*; Morgan Kaufmann Publishers Inc.: Burlington, MA, USA, 1990; pp. 396–404.
11. Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **2014**, *15*, 1929–1958.
12. Nair, V.; Hinton, G.E. Rectified linear units improve restricted boltzmann machines. In Proceedings of the 27th International Conference on Machine Learning (ICML-10), Toronto, ON, Canada, 21–24 June 2010; pp. 807–814.
13. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25*; Pereira, F., Burges, C.J.C., Bottou, L., Weinberger, K.Q., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2012; pp. 1097–1105.
14. Krizhevsky, A. Learning Multiple Layers of Features from Tiny Images. Available online: <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf> (accessed on 15 December 2017).
15. Sonka, M.; Hlavac, V.; Boyle, R. *Image Processing, Analysis, and Machine Vision*; Cengage Learning: Boston, MA, USA, 2014.
16. Gonzalez, R.; Wintz, P. *Digital Image Processing*; Addison-Wesley Pub. Co.: Boston, MA, USA, 1977.
17. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. *arXiv* **2015**, arXiv:1512.03385.
18. LeCun, Y.; Cortes, C. MNIST Handwritten Digit Database. Available online: <http://yann.lecun.com/exdb/mnist/> (accessed on 15 December 2017).
19. Pratt, H.; Williams, B.M.; Ku, J.; Coenen, F.; Zheng, Y. Automatic Detection and Identification of Retinal Vessel Junctions in Colour Fundus Photography. In Proceedings of the 21st Annual Conference on Medical Image Understanding and Analysis (MIUA 2017), Edinburgh, UK, 11–13 July 2017; Springer International Publishing: Cham, Switzerland, 2017; pp. 27–37.
20. Abbasi-Sureshjani, S.; Smit-Ockeloen, I.; Zhang, J.; Ter Haar Romeny, B. Biologically-inspired supervised vasculature segmentation in SLO retinal fundus images. In Proceedings of the 12th International Conference on Image Analysis and Recognition (ICIAR 2015), Niagara Falls, ON, Canada, 22–24 July 2015; Springer International Publishing: Cham, Switzerland, 2015.
21. Chollet, F. keras. Available online: <https://github.com/fchollet/keras> (accessed on 15 December 2017).
22. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. Available online: <https://arxiv.org/pdf/1412.6980.pdf> (accessed on 15 December 2017).

