



Heriot-Watt University
Research Gateway

The topology-overlap trade-off in retinal arteriole-venule segmentation

Citation for published version:

Juanco Muller, AV, Mota, JFC, Goatman, KA & Hoogendoorn, C 2023, The topology-overlap trade-off in retinal arteriole-venule segmentation. in O Colliot & I Isgum (eds), *Medical Imaging 2023: Image Processing.*, 124642I, Proceedings of SPIE, vol. 12464, SPIE, SPIE Medical Imaging 2023, San Diego, California, United States, 19/02/23. <https://doi.org/10.1117/12.2654014>

Digital Object Identifier (DOI):

[10.1117/12.2654014](https://doi.org/10.1117/12.2654014)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Peer reviewed version

Published In:

Medical Imaging 2023: Image Processing

Publisher Rights Statement:

©2023 Society of Photo-Optical Instrumentation Engineers (SPIE). Downloading of the abstract is permitted for personal use only.

Proceedings Volume 12464, Medical Imaging 2023: Image Processing; 124642I (2023)
<https://doi.org/10.1117/12.2654014>

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Heriot-Watt University
Research Gateway

The Topology-Overlap Trade-off in Retinal Arteriole-Venule Segmentation

Citation for published version:

Juanco Muller, AV, Mota, JFC, Goatman, KA & Hoogendoorn, C 2023, The Topology-Overlap Trade-off in Retinal Arteriole-Venule Segmentation. in *SPIE Medical Imaging*.

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Published In:

SPIE Medical Imaging

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

The Topology-Overlap Trade-off in Retinal Arteriole-Venule Segmentation

Ángel Víctor Juanco-Müller^{a,b}, João F.C. Mota^b, Keith A. Goatman^a, and Corné Hoogendoorn^a

^aCanon Medical Research Europe LTD, 2 Anderson Pl, EH6 5NP, Edinburgh, United Kingdom

^bHeriot-Watt University, Campus The Avenue, EH14 4AS, Edinburgh, United Kingdom

ABSTRACT

Retinal fundus images can be an invaluable diagnosis tool for screening epidemic diseases like hypertension or diabetes. And they become especially useful when the arterioles and venules they depict are clearly identified and annotated. However, manual annotation of these vessels is extremely time demanding and taxing, which calls for automatic segmentation. Although convolutional neural networks can achieve high overlap between predictions and expert annotations, they often fail to produce topologically correct predictions of tubular structures. This situation is exacerbated by the bifurcation versus crossing ambiguity which causes classification mistakes. This paper shows that including a topology preserving term in the loss function improves the continuity of the segmented vessels, although at the expense of artery-vein misclassification and overall lower overlap metrics. However, we show that by including an orientation score guided convolutional module, based on the anisotropic single sided cake wavelet, we reduce such misclassification and further increase the topology correctness of the results. We evaluate our model on public datasets with conveniently chosen metrics to assess both overlap and topology correctness, showing that our model is able to produce results on par with state-of-the-art from the point of view of overlap, while increasing topological accuracy.

Keywords: Retinal fundus images, convolutional neural networks, topology loss function, orientation scores

1. INTRODUCTION

The fundus of the eye provides a window to the human microvasculature and can be imaged easily from 2D colour photographs. Whenever an accurate arteriole-venule (AV) segmentation of the retinal fundus image is available, it becomes possible to compute important biomarkers, like the arteriole to venule ratio, which can be used to diagnose conditions such as hypertension and diabetes.^{1,2} However, the required blood vessel delineation and classification are intensive and laborious tasks, which may create bottlenecks in clinicians workloads. As a result, automatic arteriole-venule (AV) segmentation models are of great interest.

A variety of automatic image segmentation algorithms have been proposed, including convolutional neural networks (CNN) which achieve state-of-the-art results in segmentation tasks.³ However, such networks are often assessed using overlap-specific metrics such as the Dice and Jaccard scores,⁴ which do not take into account desirable features for tree-like tubular segmentation. Therefore, we consider relevant measures, which we term topology metrics,^{5,6} and use them to evaluate the proposed solution.

To include tubular priors in a convolutional neural network, we use a topology loss and an orientation score guided module. Whereas the former focuses the loss on foreground pixels that are closer to the centerline rather than to the boundaries, which reduces discontinuities in the predicted vessels, the latter leverages orientation-sensitive filters implemented with cake wavelets,⁷⁻⁹ which help to disambiguate crossing regions. We evaluate our model in the DRIVE-RITE^{10,11} and LES-AV¹² public datasets and find that it performs competitively with the state-of-the-art of AV segmentation¹³⁻¹⁵ in terms of overlap metrics while consistently achieving superior topology metrics.

The paper is structured as follows. We review related work in §2, describe the network architecture, loss function and experimental details in §3, present the results in §4, and discuss and conclude the paper in §5.

Further author information: (Send correspondence to Ángel Víctor Juanco-Müller)
E-mail: victor.juancomuller@mre.medical.canon

2. RELATED WORK

Researchers have tackled retinal AV segmentation with a variety of methods, ranging from specifically crafted traditional algorithms to powerful although less interpretable deep learning models. The problem, however, remains challenging. Because fundus images provide a 2D projection of the 3D arteriole and venule trees,¹⁷ differentiating between bifurcations and intersections in these images becomes an NP-hard problem.¹⁸

Traditional methods for AV segmentation generally consist of two stages: vessel segmentation and classification. For the former task, filters suited for tubular detection, such as the *Frangi* and *Hessian* kernels,^{19,20} are used together with morphological operations²¹ to obtain vessel segmentations. For the latter, machine learning classifiers are employed to classify vessel segments as arterioles or venules.^{22,23}

Because oxygenated haemoglobin absorbs less light between the 600 and 800 nm wavelength window,²⁴ arteries reflect more light and hence are brighter than veins. Thus, in theory colour intensity could be used for classification. However, inhomogeneous illumination conditions are common in retinal fundus images, requiring illumination preprocessing stages.²⁵ Furthermore, intensities measured at the smaller vessel segments may be noisier and less useful for classification.²⁶

Some approaches connect the different segments to build a graph where the nodes are called *junctions points*, which can be of three types: endpoints, bifurcations and crossings. The task then consists of identifying subgraphs that do not contain crossing nodes and correspond to the true vessel trees. Dashtbozorg et al.²⁷ approaches this task with a graph analysis technique, and Joshi et al.²⁶ with a Dijkstra graph search. However, these approaches are limited by the accuracy of the vessel segmentation used to build the graph.

To mitigate this issue, some methods aim to extract and classify the junction points directly from the images. For this task, anisotropic wavelets, such as the single sided cake wavelet,⁷ have been used to extract a set of orientation responses. These features, when combined with other vessel segmentation methods, have proved to be helpful for junction point classification.^{8,9}

Nevertheless, the presence of pathology can alter the eye appearance in the fundus images^{1,2} and therefore decrease the performance of the above mentioned algorithms. For example, hypertension increases vessel sinuosity and diabetic retinopathy produces lesions that can change colour and intensity of the vessels.²⁸ These issues may be rectified by deep learning models trained in sufficiently large datasets with enough variability of cases. For this reason, solutions based on convolutional neural networks (CNNs)^{3,29,30} have become popular among researchers.

To learn the topology of the vessels, Hemelings et al.¹³ increase the receptive field in CNNs with dilated convolutions. To capture both thicker and thinner vessels, Kang et al.³¹ combine different kernel sizes.³² Galdran et al.¹⁴ consider both crossings and uncertain vessels to be of the same class. More conveniently, Morano et al.¹⁶ build a ground truth image whose channels correspond to continuous arteriole, venule and vessels, and minimize the binary cross entropy separately for each class, which demonstrated improved performance, and it is the strategy we adopt here. Recently, Luo et al.¹⁵ proposed an segmentation network followed by a conditional generative adversarial model for AV refinement, which also inspires the two-stage nature of the method hereby presented. Nevertheless, all these methods still fail to preserve the continuity of segmented arterioles and veins.

3. MATERIALS & METHODS

Here we describe the preprocessing stage in §3.1, the proposed model architecture in §3.2, the loss function with topology terms in §3.3 and give details of the experimental set-up in §3.4.

3.1 Preprocessing

Before feeding the images to the proposed model, we perform two preprocessing stages. First, to mitigate the heterogeneous illumination of fundus images, we apply an intensity homogenization procedure²⁵ based in the dark channel prior.³⁷ Next, we enhance the vessels by subtracting to the image its Gaussian filter response. This has the effect of a high pass filter and separates the modes of foreground and background intensity distributions.¹⁶ Figure 1 shows the different stages for a sample of the RITE dataset^{10,11} affected by inhomogeneous illumination.



Figure 1. Preprocessing of sample *37_training* DRIVE-RITE dataset.^{10,11} Left: original image. Middle: image after illumination correction.^{25,37} Right: image after vessel enhancement.¹⁶

3.2 Network architecture

The proposed neural network architecture for AV segmentation, depicted in the top of Fig. 2, broadly consists of two feature pyramid networks, explained in §3.2.1, and the cake wavelet module, described in §3.2.3. The building blocks of the feature pyramid networks are residual multi-dilation convolutional layers, as explained in §3.2.2. Two fusing blocks, presented in §3.2.4, process the output features of the feature pyramid network and cake wavelet module to produce the output segmentation results. The first feature pyramid block is responsible for segmenting the vessels, whereas the second one, conditioned on the vessel probability and input image and leveraging the orientation sensitivity of the cake wavelets, produces the AV segmentation.

3.2.1 Feature pyramid network

The feature pyramid network has three levels of depth. At the first level, the output feature map contains 16 channels, and this number is doubled at every block. The resulting feature maps are both downscaled by a factor of 2 with max-pooling for downstream processing in the feature pyramid network and resized to the original image size for input to the fusing block.

3.2.2 Residual multi-dilation block

The multi-dilation block applies convolutions of kernel size 5x5 with dilation values of 1, 2 and 3 in parallel. The output responses are then concatenated and processed by a 1x1 convolution and instance normalization layers³³ to produce the output feature map, as shown in the middle left of Fig. 2. This operation is applied three consecutive times with a residual connection³⁴ between the second and the third output, as depicted in the middle right of Fig. 2.

3.2.3 Cake wavelet module

The single-sided cake wavelets⁷ are defined as the inverse 2D Fourier transform, \mathbf{x} , of an oriented shape in frequency domain, $\tilde{\mathbf{x}}$. With an appropriate window size, the real and imaginary parts of these filters approximate an *orientation detector* perpendicular to the orientation in frequency domain (see Fig. 3, left).

The cake wavelet module first maps the RGB fundus image to grayscale with a 1x1 convolution and then convolves it with 7x7 convolutions initialized with the negative real part of the single side cake wavelets filters for 24 values of an orientation angle $\theta \in \{0, 2\pi\}$. An schematic of the block is shown in the bottom left of Fig. 2.

Although we initialize the cake wavelet module with the cake wavelet filters, we let the network learn them during training. It turns out that after training the values of the filters remain practically unchanged, which shows the suitability of the cake wavelet for the task at hand (see Fig. 3, right).

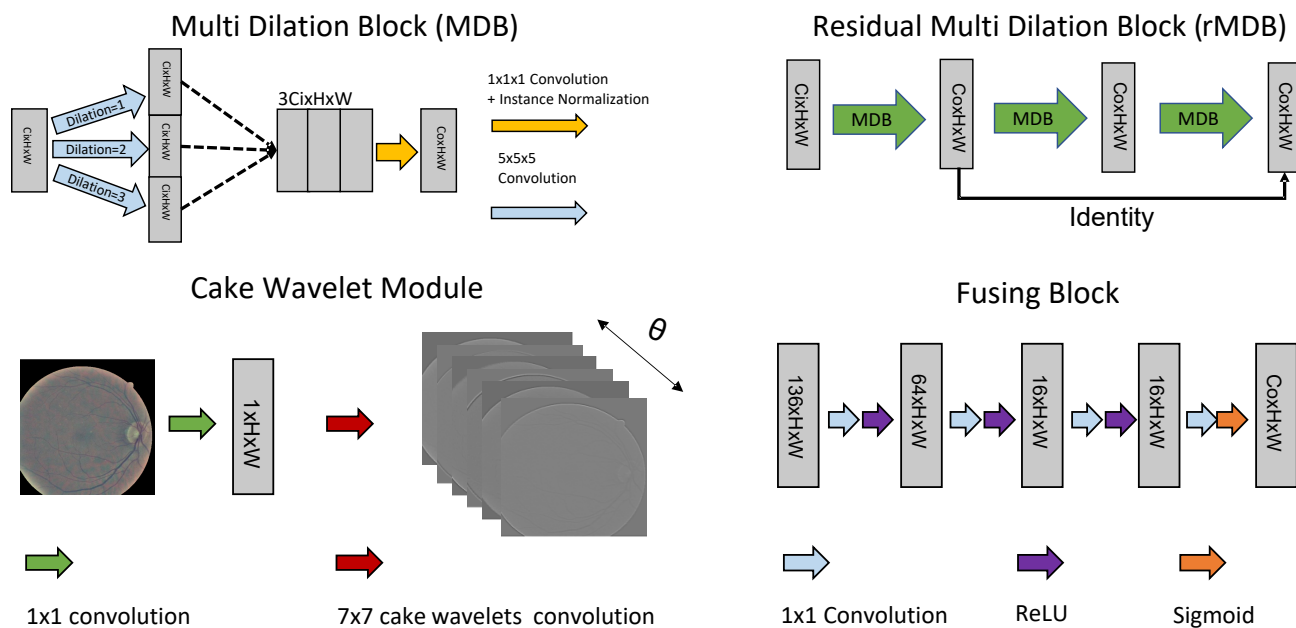
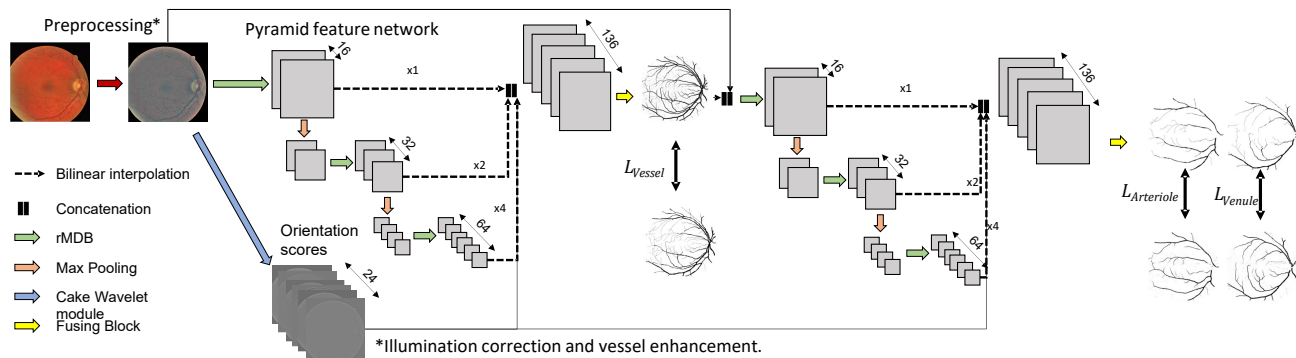


Figure 2. Top: proposed pipeline; from left for right: preprocessing, neural network processing and loss calculation. Middle, from left to right: multi-dilation block (MBD) which combines 5x5 convolutions at different dilation and its residual implementation (rMDB). Bottom, from left to right: Cake wavelet module, with a 7x7 convolution initialized with the negative real part of cake wavelets, and the fusing block, a combination of 1x1 convolutions and ReLU activations except at the final layer, where sigmoid is used.




Figure 3. Animation of the cake wavelets for varying θ . From left to right: oriented shape in frequency domain $\tilde{\mathbf{x}}$, real and imaginary part of the inverse Fourier transform \mathbf{x} , negative real part of \mathbf{x} before and after training of the CNN model ϕ .

3.2.4 Fusing block

The fusing block takes the concatenation of outputs at the feature pyramid module at different resolution levels and the orientation scores produced by the cake wavelet module. These are all concatenated resulting in a feature map with 136 channels. A series of 1x1 convolution layers and ReLU activations reduce the channel dimension from 136 to 64, from 64 to 16, and finally from 16 to the output number of segmentation channels.

3.3 Loss Function

We optimize a combination of topology and overlap losses. We set the Dice loss function L_{Dice} as the overlap term. To improve the connectivity of the predicted vessels, a topology term enhances the vessel centerline, that is the inner continuous line equidistant to both vessel walls. We adopt the centerline-Dice $L_{\text{cl-Dice}}$ proposed by Suprosanna and Shit,³⁶ described in §3.3.1.

Furthermore, we include *mixed terms*, which are the Mean Square Error L_{MSE} and Binary Cross Entropy L_{BCE} losses weighed by a centerline cost-map α according to the policy described in §3.3.2. Thus, the network is trained to minimize the following loss

$$L_{\text{class}} = \underbrace{\lambda_1 L_{\text{Dice}}}_{\text{Overlap term}} + \underbrace{\lambda_2 L_{\text{cl-Dice}}}_{\text{Topology term}} + \underbrace{\alpha (\lambda_3 L_{\text{MSE}} + \lambda_4 L_{\text{BCE}})}_{\text{Mixed terms}}, \quad (1)$$

where $\text{class} \in \{\text{arteriole, venule, vessels}\}$, and $\lambda_1, \lambda_2, \lambda_3$ and λ_4 are constants experimentally set to 1, 0.5, 0.5 0.5 respectively.

3.3.1 Centerline Dice

The centerline Dice³⁶ applies max and min pool functions to a vessel likelihood map for a number of iterations k to retrieve the centerline or skeleton. Because min and max pool are differentiable functions, the resulting skeleton probability can be used to compute the Dice loss with respect to the skeleton ground truth and backpropagate gradients to directly improve the structural connectivity of the outputs. Fig. 4 left, shows how the number of iterations k relates with the quality of the centerline; in our work setting $k = 5$ was sufficient to achieve good results without considerably increasing complexity.

3.3.2 Centerline Cost Map

To balance the scale of the gradients a *cost-map* α is used to weight pixels in the centerline more than those on the vessel walls. The cost-map α is defined as the geodesic distance transform with respect to the ground truth centerline restricted to the ground truth vessels (Fig. 4, right). However, to allow the network first learn the overall vessel shape, and then correct for connectivity mistakes, α weights all pixels equally until the vessel predictions reach a minimum Dice score (~ 0.6), and only after that α is set to the described centerline distance transform.

3.4 Experimental Set-up

This section explains our implementation details, starting with data augmentation in §3.4.1, and then presenting the used hardware, software and other training details in §3.4.2.

3.4.1 Data Augmentation

We trained our network on 17 images from the DRIVE-RITE training set^{10,11} and withheld the three remaining samples to validate the model after each epoch. To enlarge the number of training samples we applied a variety of augmentation techniques: rotations, flips, intensity shifts and elastic deformations were used to build eight augmented patches from each training sample. To better capture vessel structure, we set a large patch size (512×512 ; the full images are 584×565). At inference we used a 512×512 sliding window with a 10% overlap.




Figure 4. Left: animation of the iterative process of soft skeletonization. Setting $k=5$ was sufficient to achieve good results. Right: centerline geodesic distance transform to weigh foreground pixels in the loss computation.

3.4.2 Software, Hardware and Training details

The preprocessing stages were implemented with the Scikit-Image library.³⁸ For the neural networks, we used the Pytorch Deep Learning framework³⁹ with the ADAM optimizer.⁴² The MONAI library⁴⁰ was used for data augmentation and loading. For the geodesic distance transforms, we used the scikit-fmm package.⁴¹

We tracked a combination of two overlap (Dice score and precision) and two topology (tree detection rate and branch detection rate) metrics.⁶ We assumed convergence whenever the tracked metric did not improve for more than 10 epochs.

The experiments were run using a NVIDIA Tesla V100 with 32 GB of RAM. The average memory and time by training epoch was 5 GB and 3 minutes with batch size of 1. The learning rate, initially set to 0.001, was reduced by a factor of 0.1 whenever this metric did not improve for 5 epochs. Total training was completed between 2 and 3 hours.

4. RESULTS

This section presents the AV segmentation results. We give a first evaluation with the receiving operator characteristic (ROC) curve in §4.1. We then introduce the topology sensitive metrics in §4.2. Next, we present the results in the DRIVE-RITE^{10,11} dataset in §4.3, and finally the results on the LES-AV dataset¹² in §4.4.

4.1 ROC Analysis

Fig. 5 presents the ROC of the baseline model, and the two proposed modifications, computed on the 20 test images of the DRIVE-RITE dataset.^{10,11} It shows that all three models achieve high areas under the curve (AuCs), with a slight improvement for the model with both topology loss and Cake wavelets. Given the high AuC for all models, we decided to use a threshold of 0.5 to binarize the model outputs.




Figure 5. ROC curves computed on the 20 test images of the DRIVE-RITE dataset^{10,11} for the three proposed models. As it can be seen, the version with both topology loss and cake wavelets results in the curve enclosing the highest area.

4.2 Overlap, Mixed & Topology Metrics

To assess both overlap and topology correctness we devise specific metrics and organise them in three groups: *overlap*, *mixed*, and *topology* ones.

The *overlap* and *mixed* metrics are derived from the evaluation code released by Hemelings et al.,¹³ and they calculate the F1 and accuracy of the AV classification on all pixels (*overlap*) and only on the ground truth centerline pixels (*mixed* metrics). The *topology* metrics are the vessel detection rate, also from Hemelings et al.,¹³ and the tree length and branch detection rates, from the Airway Tree Modelling 2022 (ATM22) challenge.⁶

The vessel detection rate is the prediction accuracy computed only on ground truth vessel pixels. The tree length detection rate is the sensitivity of the prediction with respect to the ground truth centerline. Finally, the branch detection rate uses an image where ground truth vessel segments are uniquely labelled to determine how many of them are present in the predicted segmentation.

Both the *topology* and *mixed* metrics are more difficult to optimize because they provide local measurements, constrained to pixels in certain regions of the images, and therefore are more sensitive to the predicted vessel topology than overlap measurements evaluated in all pixels.

4.3 AV Segmentation Results on DRIVE-RITE

We compared the baseline model against the version with topology loss and the one with topology loss and Cake wavelet module. For each combination, we trained 5 models and used an ensemble of them to produce the final output. Table 1 reports results using the obtained segmentations and the masks shared by Hemelings et al.¹³

Although our model performs slightly worse than Hemelings et al.¹³ at the all-pixel level, it outperforms their method at the centerline level and in terms of the topology metrics. The lower overlap quality is a side effect of optimizing the models for topology correctness, which is mitigated with the cake wavelet module. Further design strategies are needed to balance the topology-overlap trade-off.

Finally, to give a qualitative interpretation of the results, Fig. 6 illustrates the effect of the proposed modifications to the baseline model. As it can be seen, the model trained with the topology loss function produces continuous vessels, however at the cost of misclassifying arterioles and venules pixels within them. Including the Cake Wavelet module in addition to the topology loss mitigates this issue.

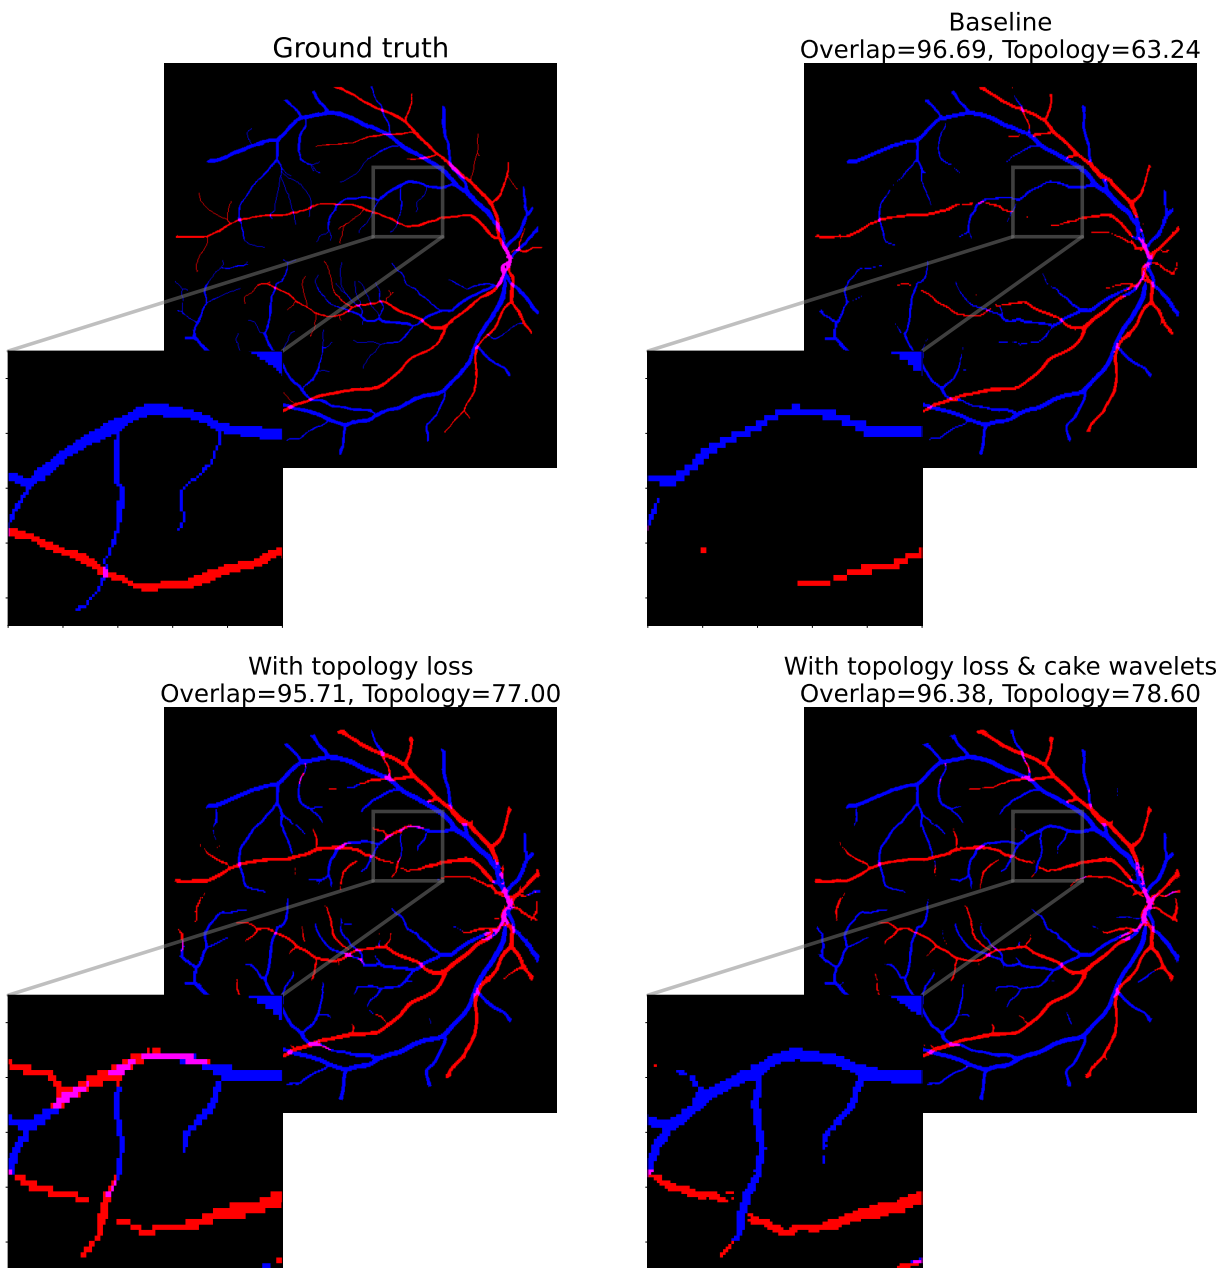


Figure 6. Top left: Ground truth of sample *test07* of the DRIVE-RITE dataset.^{10,11} Top right: the prediction of our baseline model, which achieves high overlap at the cost of continuity mistakes. Bottom left: the result of model with topology loss, which produces continuous vessels although with misclassified regions. Bottom right: the output of the model trained with both topology loss and orientation scores, which produces continuous and correctly classified vessels. Overlap is the average of F1 and accuracy as in,¹³ and topology is the average of tree, branch⁶ and vessel detection rates.¹³

Table 1. AV classification results in the RITE test set.^{10,11} Although our method achieves slightly worse results than Hemelings et al.¹³ at the overlap level, it performs best in the topology and mixed metrics. Best figures are in bold.

	Overlap metrics (evaluated on all pixels)		Mixed Metrics (evaluated on centerline pixels)		Topology metrics (detection rates)		
	F1	Accuracy	F1	Accuracy	Branches	Tree length	Vessel
Baseline	95.98	95.85	55.88	43.73	68.75	56.73	66.63
Hemelings et al. ¹³	96.78	96.71	66.65	50.97	76.75	66.61	73.06
Baseline + topology loss	94.97	95.31	68.99	58.27	76.32	72.13	78.14
Baseline + topology loss + cake wavelets	95.66	95.87	69.85	59.96	76.84	72.93	77.85

Table 2. Results in the LES-AV dataset¹² of the model trained on RITE^{10,11} compared to other researcher’s methods. Best figures are in bold.

Method	Accuracy	F1	Specificity
Galdran et al. (2019) ¹⁴	86.00	86.00	85.00
Luo et al. (2022) ¹⁵	95.10	96.10	94.10
Ours	96.12	96.40	97.20

4.4 AV Segmentation Results on LES-AV

To explore the generalizability of the proposed model to unseen images we apply it to the LES-AV dataset,¹² which contains fundus images and AV segmentations of size 1620×1444 . We resized the samples in LES-AV¹² with bilinear interpolation to the patch size used for training, that is 512×512 , and retrieve the binary reference segmentations by thresholding at 0.5.

Table 2 presents AV classification results compared to other state of the art approaches in the LES-AV dataset.¹² Because topology metrics are not reported in these papers, we limit our assessment to overlap (pixel-wise) measures. We again used the code released by Hemelings et al.¹³ to compute the metrics, and reported results in other authors’ papers. In spite of being trained with only 17 images from the RITE dataset,^{10,11} our method performs competitively with the LES-AV dataset.¹²

5. DISCUSSION & CONCLUSION

We have proposed a method to tackle the overlap-topology trade-off in arteriole-venule segmentation in fundus images. Our algorithm, however, is general: it is applicable to any task involving segmentation of complex tubular structures. For this reason, we adopted metrics used for pulmonary airway segmentation,⁶ and used them to validate the effect of the topology loss terms in our loss function.

We found that only using topology loss increases the continuity of the segmented vessels but results in poorer overlap-based quality, which translates into arteriole-venule classification mistakes. This issue was addressed by incorporating an orientation-score guided module using cake wavelets, a type of orientation sensitive filter which was previously used for junction point classification.^{8,9}

In conclusion, we novelly incorporate topology preserving losses and orientation-scores in a CNN-based vision model to obtain semantically accurate retinal AV segmentations. Both improving the cake-wavelet module design and extending the method to other modalities such as 3D computed tomography angiography are important directions for future work.

REFERENCES

- [1] Stokoe, N. L., and R. W. Turner. "Normal retinal vascular pattern. Arteriovenous ratio as a measure of arterial calibre." *The British journal of ophthalmology* 50.1 (1966): 21.
- [2] Vázquez, S. G., et al. "Improving retinal artery and vein classification by means of a minimal path approach." *Machine vision and applications* 24.5 (2013): 919-930.
- [3] Long, Jonathan, Evan Shelhamer, and Trevor Darrell. "Fully convolutional networks for semantic segmentation." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015.
- [4] Taha, Abdel Aziz, and Allan Hanbury. "Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool." *BMC medical imaging* 15.1 (2015): 1-28.
- [5] Rubenhx/AV-segmentation. GitHub. Available at: <https://github.com/rubenhx/av-segmentation> (Accessed: December 27, 2022).
- [6] ATM-22-related-work [MICCAI, 2022, challenge]: Airway tree modeling (ATM'22) related work collections" Available at: <https://github.com/Puzzled-Hui/ATM-22-Related-Work> (Accessed: December 27, 2022).
- [7] Duits, R., Duits, M., van Almsick, M. et al. Invertible orientation scores as an application of generalized wavelet theory. *Pattern Recognit. Image Anal.* 17, 42–75 (2007). <https://doi.org/10.1134/S1054661807010063>
- [8] Hurtado Martí, Feliciano. "Implementación de un banco de filtros de Cake Wavelets para la estimación de la orientación múltiple y su uso en el seguimiento de la vasculatura en imágenes retinográficas." (2016).
- [9] Abbasi-Sureshjani, Samaneh, et al. "Automatic detection of vascular bifurcations and crossings in retinal images using orientation scores." 2016 IEEE 13th international symposium on biomedical imaging (ISBI). IEEE, 2016.
- [10] Staal, J., Abràmoff, M., Niemeijer, M., Viergever, M., van Ginneken, B., 2004. Ridge based vessel segmentation in color images of the retina. *IEEE Trans. Med. Imaging* 23, 501–509
- [11] Qureshi, Touseef Ahmad, et al. "A manually-labeled, artery/vein classified benchmark for the DRIVE dataset." *Proceedings of the 26th IEEE international symposium on computer-based medical systems*. IEEE, 2013.
- [12] Orlando, José Ignacio, et al. "Towards a glaucoma risk index based on simulated hemodynamics from fundus images." *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, Cham, 2018.
- [13] Hemelings, Ruben, et al. "Artery-vein segmentation in fundus images using a fully convolutional network." *Computerized Medical Imaging and Graphics* 76 (2019): 101636.
- [14] Galdran, Adrian, et al. "Uncertainty-aware artery/vein classification on retinal images." 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). IEEE, 2019.
- [15] Luo, Shichen, et al. "Two-Stage Topological Refinement Network for Retinal Artery/Vein Classification." 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI). IEEE, 2022.
- [16] Morano, José, et al. "Simultaneous segmentation and classification of the retinal arteries and veins from color fundus images." *Artificial Intelligence in Medicine* 118 (2021): 102116.
- [17] Estrada, Rolando, et al. "Retinal artery-vein classification via topology estimation." *IEEE transactions on medical imaging* 34.12 (2015): 2518-2534.
- [18] Estrada, Rolando, et al. "Tree topology estimation." *IEEE transactions on pattern analysis and machine intelligence* 37.8 (2014): 1688-1701.
- [19] Frangi, A. F., Niessen, W. J., Vincken, K. L., & Viergever, M. A. (1998,). Multiscale vessel enhancement filtering. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 130-137). Springer Berlin Heidelberg
- [20] Sato, Yoshinobu, et al. "Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images." *Medical image analysis* 2.2 (1998): 143-168.
- [21] A. Mendonça and A. Campilho, "Segmentation of retinal blood vessels by combining the detection of centerlines and morphological reconstruction," *IEEE Trans. Med. Imag.*, vol. 25, no. 9, pp. 1200–1213, Sep. 2006.
- [22] Akbar, Shahzad, et al. "Decision support system for detection of hypertensive retinopathy using arteriovenous ratio." *Artificial intelligence in medicine* 90 (2018): 15-24.

- [23] M. M. Fraz, A. R. Rudnicka, C. G. Owen, D. P. Strachan and S. A. Barman, "Automated arteriole and venule recognition in retinal images using Ensemble classification," 2014 International Conference on Computer Vision Theory and Applications (VISAPP), 2014, pp. 194-202.
- [24] Huang, Fan, Behdad Dashtbozorg, and Bart M. Romeny. "Artery/vein classification using reflection features in retina fundus images." *Machine Vision and Applications* 29.1 (2018): 23-34.
- [25] A. Galdran, A. Alvarez-Gila, A. Bria, J. VazquezCorral, and M. Bertalmio. On the Duality Between Retinex and Image Dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [26] Joshi, Vinayak S., et al. "Automated method for identification and artery-venous classification of vessel trees in retinal vessel networks." *PloS one* 9.2 (2014): e88061.
- [27] Dashtbozorg, Behdad, Ana Maria Mendonça, and Aurélio Campilho. "An automatic graph-based approach for artery/vein classification in retinal images." *IEEE Transactions on Image Processing* 23.3 (2013): 1073-1083.
- [28] Huadeng Wang, Guang Xu, Xipeng Pan, Zhenbing Liu, Ningning Tang, Rushi Lan, Xiaonan Luo. "Attention-inception-based U-Net for retinal vessel segmentation with advanced residual." *Computers & Electrical Engineering* 98 (2022): 107670.
- [29] Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." *International Conference on Medical image computing and computer-assisted intervention*. Springer, Cham, 2015. **2**(5), 99–110 (2016)
- [30] Isensee, Fabian, et al. "nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation." *Nature methods* 18.2 (2021): 203-211.
- [31] Kang, Hong, et al. "AVNet: A retinal artery/vein classification network with category-attention weighted fusion." *Computer Methods and Programs in Biomedicine* 195 (2020): 105629.
- [32] Szegedy, Christian, et al. "Going deeper with convolutions." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015.
- [33] Ulyanov, D., et al. (2016). Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*.
- [34] He, Kaiming, et al. "Deep residual learning for image recognition." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016.
- [35] Milletari, Fausto, Nassir Navab, and Seyed-Ahmad Ahmadi. "V-net: Fully convolutional neural networks for volumetric medical image segmentation." 2016 fourth international conference on 3D vision (3DV). IEEE, 2016.
- [36] Suprosanna Shit, Johannes C. Paetzold, Anjany Sekuboyina, Ivan Ezhov, Alexander Unger, Andrey Zhylyka, Josien P. W. Pluim, Ulrich Bauer, Bjoern H. Menze; *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 16560-16569
- [37] Kaiming He, Jian Sun and Xiaoou Tang, "Single image haze removal using dark channel prior," 2009 IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 1956-1963, doi: 10.1109/CVPR.2009.5206515.
- [38] Stéfan van der Walt, Johannes L. Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D. Warner, Neil Yager, Emmanuelle Gouillart, Tony Yu and the scikit-image contributors. *scikit-image: Image processing in Python*. *PeerJ* 2:e453 (2014) <https://doi.org/10.7717/peerj.453>
- [39] Paszke, Adam, et al. "Pytorch: An imperative style, high-performance deep learning library." *Advances in neural information processing systems* 32 (2019).
- [40] The MONAI Consortium. Project MONAI. Zenodo. <http://doi.org/10.5281/zenodo.4323059> (Accessed: December 27, 2022).
- [41] Scikit-Fmm: the fast marching method for Python. GitHub. Available at: <https://github.com/scikit-fmm/scikit-fmm> (Accessed: December 27, 2022).
- [42] Kingma, Diederik P., and Jimmy Ba. "Adam: A method for stochastic optimization." *arXiv preprint arXiv:1412.6980* (2014).