Computers in Biology and Medicine 195 (2025) 110516

Contents lists available at ScienceDirect

Computers in Biology and Medicine

journal homepage: www.elsevier.com/locate/compbiomed

Towards a comprehensive characterization of arteries and veins in retinal imaging

Paolo Andreini, Simone Bonechi *

Department of Information Engineering and Mathematics (DIISM), University of Siena, Via Roma 56, I–53100, Siena, Italy

ARTICLE INFO

Keywords:
Medical Image Segmentation and localization
Arteries and veins segmentation
Vessel reconstruction and characterization
Retinal fundus analysis

ABSTRACT

Retinal fundus imaging is crucial for diagnosing and monitoring eye diseases, which are often linked to systemic health conditions such as diabetes and hypertension. Current deep learning techniques often narrowly focus on segmenting retinal blood vessels, lacking a more comprehensive analysis and characterization of the retinal vascular system. This study fills this gap by proposing a novel, integrated approach that leverages multiple stages to accurately determine vessel paths and extract informative features from them. The segmentation of veins and arteries, achieved through a deep semantic segmentation network, is used by a newly designed algorithm to reconstruct individual vessel paths. The reconstruction process begins at the optic disc, identified by a localization network, and uses a recurrent neural network to predict the vessel paths at various junctions. The different stages of the proposed approach are validated both qualitatively and quantitatively, demonstrating robust performance. The proposed approach enables the extraction of critical features at the individual vessel level, such as vessel tortuosity and diameter. This work lays the foundation for a comprehensive retinal image evaluation, going beyond isolated tasks like vessel segmentation, with significant potential for clinical diagnosis.

1. Introduction

Retinal fundus imaging plays a crucial role in ophthalmology, enabling visualization of the eye's posterior segment and aiding in the diagnosis and monitoring of various eye diseases [1,2]. A key component of retinal image analysis is the segmentation and characterization of blood vessels. Vessels provide valuable information about the health of the eye, with their structure and morphology being associated with various pathologies [3,4] Moreover, the retinal microvasculature is the only part of human circulation that can be directly and non-invasively visualized in vivo. As a result, analyzing vascular structures in retinal images plays an important role in detecting systemic disorders such as diabetes and hypertension [5]. Early and accurate detection of retinal vascular changes is critical to prevent vision loss in diabetic retinopathy, retinopathy of prematurity, and other cardiovascular complications [6]. Deep learning has significantly advanced retinal image analysis [7], with many methods focusing on vessel segmentation [8,9] and some specifically aiming to differentiate between veins and arteries [10]. Few approaches attempt to go beyond vessel segmentation to characterize vessels by extracting geometric features. However, these methods often do not distinguish between veins and arteries, require manual extraction of the vessels, or consider only sub-parts, relying on inaccurate traditional computer vision techniques for vessel segmentation [11–13]. The lack of an effective, comprehensive method for both vessel segmentation and characterization motivates this work. This paper proposes a novel approach that goes beyond segmentation, aiming to automatically extract and reconstruct the retinal vessels, providing a more comprehensive analysis for improved clinical decision support. Our system leverages a multi-stage approach (see Fig. 1). In an initial phase, we employed a YOLO (You Only Look Once) detection network [14] to localize the optic disc and a SegNeXt segmentation network [15] to segment arteries and veins. Next, we use skeletonization on the segmentation maps to obtain an estimate of the vessel structure, and the main vessels are reconstructed from their start points (seeds) to their end points. The seeds are determined by the intersection of the skeleton with the optic disc. From these seeds, the reconstruction method is applied to reconstruct the vessels, employing a Recurrent Neural Network (RNN) to predict their continuation at the junctions. Specifically, two models are compared: a Long Short-Term Memory (LSTM) network [16] and a bidirectional LSTM (BiLSTM) network [17]. During reconstruction, the algorithm iteratively extends the vessel path

* Corresponding author.
E-mail address: simone.bonechi@unisi.it (S. Bonechi).

https://doi.org/10.1016/j.compbiomed.2025.110516
Received 11 February 2025; Received in revised form 13 May 2025; Accepted 1 June 2025
Available online 23 June 2025

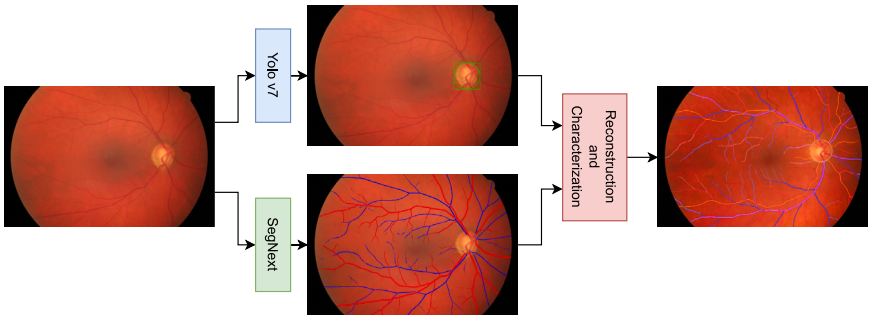

Fig. 1. Vessel reconstruction and characterization scheme: Veins and arteries segmentation is provided by the SegNext model while the Yolov7 is used to localize the optic disc.

from the seed point, following the skeleton, employing the RNN to predict the most probable continuation based on the existing vessel segment. Any unselected branch points are stored as additional seed points, allowing the method to iteratively map secondary, tertiary, and higher-order vessel branches throughout the retinal vascular network. This process enables the effective reconstruction of all the retinal vessels. From these reconstructions, we extract geometrical features such as tortuosity, length, and diameter, which are known indicators of various eye diseases [18]. Our comprehensive characterization of the retinal vasculature can significantly aid clinicians, potentially facilitating early disease detection, monitoring disease progression, and ultimately improving patient care.

The paper is organized as follows. The Related Works (Section 2) discusses existing retinal vessel analysis methods. The datasets and the deep learning networks employed in this study were presented in Section 3, along with the proposed system to characterize blood vessels, including optic disc localization, arteries and veins segmentation, and blood vessel reconstruction. Section 4 presents the experiments used to validate the proposed approach, whose results are discussed in Section 5. Finally, Section 6 concludes and presents some future perspectives of this work.

2. Related works

Researchers have proposed numerous methods to analyze retinal fundus images, employing different approaches for segmentation and classification. These techniques aim to automatically identify and classify various anatomical structures within the retina, such as the optic disc, blood vessels, and fovea.

Optic disc detection. A common approach for optic disc detection consists in adapting well-known machine learning models to this problem. Many methods use the UNet [19] architecture combined with specific techniques to achieve higher accuracy. For instance, in [20] probability bubbles are used to improve the results, while Meng et al. [21] utilizes a two-stage approach to address the class imbalance. The authors of Maiti et al. [22], modified the UNet employing multiple sub-networks with different models (e.g. LSTMs). Recent advancements include incorporating attention blocks and spatial pyramid networks for enhanced optic disc and cup segmentation [23].

Blood vessel segmentation. In recent years the use of deep learning techniques has played a leading role in the segmentation of blood vessels. Two major research trends stand out: the first focuses on designing specialized networks and strategies to accurately recognize vessels that may vary in size and be extremely thin, while the second emphasizes developing strategies to address the often limited training data available. To improve the segmentation performance numerous works focus on balancing the learning process to accurately capture the vessel morphology. Bridge-net [24] addresses this by combining RNNs with Convolutional Neural Networks (CNNs) and employing patch-based loss weight mapping. In this context UNet-based architectures with various modifications are prevalent. A Dense UNet designed for optimized learning with a minimal number of parameters is proposed in [25], whereas Li et al. [26] employs attention modules to capture global information and enhance feature extraction. The HAnet model uses a multi-decoder network and attention mechanism to further refine segmentation [27]. Additionally, the weighted attention gate strategy provides a method to eliminate background noise and improve segmentation accuracy [28]. Several other CNN-based methods have been proposed, including Scs-Net [29], NFN+ [30], and Sa-Nnet [31]. Instead, the scarcity of training data was addressed in [9, 32,33], all the methods employed Generative Adversarial Networks to synthesize high-quality retinal images along with the corresponding semantic label-maps that can be effectively used during the training of a segmentation network.

Artery-vein segmentation. Deep learning has significantly advanced the field of retinal Artery-Vein (A/V) segmentation and classification. A memory-efficient network for high-resolution images was proposed by Lepetit-Aimon et al. [34] but it struggles to capture the intricate branching patterns of blood vessels. To address this, cascaded Unet (CRU-Net) architectures were developed for dual-modal fundus images [35], but these methods required separate classifications for arteries and veins, increasing computational costs. In [36], the authors incorporated a category-attention mechanism for improved A/V recognition, while Hemelings et al. [37],Ma et al. [38] adopted a single UNet for simultaneous vessel segmentation and A/V classification. The authors of Girard et al. [39] proposed an encoder–decoder CNN model with pre-processing and image stacking for segmentation. Uncertainty in A/V recognition is addressed in [40] where a network is trained to identify uncertain regions alongside classification. Recognizing the inherent link between segmentation and classification, Wang et al. [41] proposed a multi-task siamese network for simultaneous learning. Another UNet-based simultaneous segmentation approach for A/V classification was proposed in [42]. The Topology and Width Aware Generative Adversarial Network (TW-GAN) was proposed in [43] to integrate topology connectivity and vessel width information into the deep learning framework for A/V classification. Interestingly, Galdran et al. [44] demonstrated that a carefully trained minimalistic UNet architecture can compete with state-of-the-art techniques. Similarly, Karlsson and Hardarson [45] proposed a CNN using serially connected UNets for simultaneous segmentation and classification. More recently, a novel semi-supervised point consistency network (SPC-Net) for A/V classification was proposed in [46]. Additionally, the RRWNet proposed in [10] allows for recursive refinement of segmentation maps and correction of classification errors. Instead of relying on deep learning models that demand large, balanced datasets, Mokan et al. [47] uses traditional machine learning with handcrafted, intensity-based features, producing competitive results compared with deep learning methods. The EAV-Net, an efficient deep learning model designed for retinal artery and vein classification that effectively addresses the challenges of vessel discontinuity and A/V confusion by incorporating a multi-scale convolution block, was introduced in [49]. Finally, [48] introduces the

Table 1
Summary of advantages and disadvantages of A/V segmentation methods.

Method	Advantages	Disadvantages
Memory-efficient FCN [34]	Low memory and computation; handles high-resolution images.	A/V classification result under state-of-the-art.
CRU-Net – Cascaded U-Net [35]	Effective dual-modal usage with high segmentation accuracy.	Dual-modal fundus cameras are not as widely used as conventional fundus cameras.
AVNet + Category-attention [36]	Enhances artery/vein discrimination using attention mechanisms.	The proposed approach relies on GoogLeNet, an outdated architecture.
Single U-Net [37,38]	End-to-end training on A/V; provides annotations for HRF.	Uses a general purpose network whose results are improved by more recent approaches.
Encoder–decoder + Graph [39]	Enforces global consistency via graph propagation.	The two-stage approach is constrained by a weak initial step.
Uncertainty-aware A/V [40]	Uses uncertainty estimation to flag ambiguous regions.	The images are processed at a resolution of 512 × 512, which may result in the loss of vessel details.
Multi-task Siamese [41]	Jointly learns segmentation and topology for high accuracy.	Sophisticated framework that may demand significant computational resources.
Simultaneous U-Net [42]	Simultaneous segmentation and classification of the retinal arteries and veins.	Lacks explicit mechanisms for consistent vessel labeling relying on relatively low level features.
TW-GAN [43]	Integrates topology and vessel width for improved coherence.	Involves complex adversarial training and requires high computational resources.
Minimalistic U-Net [44]	Achieves competitive performance with a simple, memory-efficient model.	The architecture is trained on downsampled images, which may result in the loss of vessel details.
Serial U-Nets [45]	Serially connected U-nets allow to increase the performance.	Increased model complexity and A/V mistakes are still present.
Semi-supervised SPC-Net [46]	Leverages unlabeled data to improve performance.	Sometimes the detected small blood vessels are discontinuous.
Recursive Refinement (RRWNet) [10]	Iteratively refines outputs to enhance consistency.	The approach requires fine-tuning of the iteration count for segmentation refinement.
Pixel-wise classification [47]	Demonstrates competitive results based on a standard machine learning approach.	Significant computational resources and extended execution time are required.
DBMAE-Net [48]	Good ability to accurately differentiate and segment the retinal arterial and venous vessels in complex fundus images.	The study relies on publicly available datasets, which come with a small number of images.
EAV-Net [49]	The method tackles two major challenges: vessel discontinuity (especially for small vessels) and color-based confusion between arteries and veins	While the approach helps reduce confusion between arteries and veins, it does not completely solve this problem.

dual-branch multi-scale feature adaptive extraction network (DBMAE-Net), a novel approach for simultaneous retinal arteriovenous vessel segmentation and tube diameter measurement.

The main advantages and disadvantages of all the A/V segmentation approaches are reported in Table 1.

Geometric feature extraction. Measuring the diameter of retinal blood vessels is critical for assessing eye health, and researchers have developed various automated methods to achieve this goal. In [12], a simple thresholding method is used to extract the vessels, and after a thinning step, the active contours are used to detect vessel contours that allow the calculation of vessel diameters. Similarly, in [13], the Radon Transform and Gaussian Process method are used on vessel ground truth maps to determine vessel diameter. However, this method struggles with non-Gaussian distributed vessel diameters and is computationally intensive. Another approach [50] employed the Euclidean distance transformation to calculate the distance between bright pixels within a designated region to estimate the vessel diameter. In [51], the authors focused their study on the analysis of specific vascular diameters, such as central retinal arteriolar equivalent and central retinal venular equivalent. While in [52], employing a standard image processing approach the authors detect the optic disc, differentiate between veins and arteries, and calculate both vessel diameter and arteriolar to venular diameter ratio. Other techniques include automated methods that use linear discriminant analysis [53] for diameter measurement and a difference-based two-dimensional Gaussian model [54] that achieves subpixel precision. Furthermore, a vessel tracking algorithm [55] models the grayscale distribution across the cross-section of a vessel using twin Gaussian functions to estimate the diameter. Finally, in [11], the authors develop a Gaussian rider-based approach for accurate tracking and diameter estimation of retinal blood vessels.

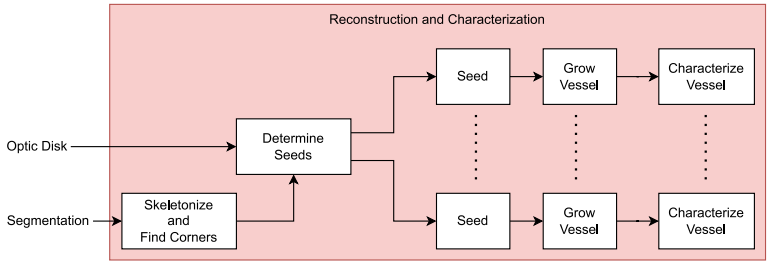

Fig. 2. An overview of the blood vessel reconstruction algorithm. Seeds are determined by intersecting the vessel skeleton and the optic disc. Each vessel is grown from the initial seed and then characterized.

3. Materials and methods

This section provides a comprehensive overview of the datasets, networks, and methodologies utilized in our study to reconstruct and characterize blood vessels in retinal fundus images. Section 3.1 describes the three datasets employed: HRF, RITE, and FIVES. We then detail the specific networks used in our approach: SegNeXt for artery/vein segmentation (Section 3.2), YOLO for optic disc localization (Section 3.3), and recurrent models (LSTM and BiLSTM) for blood vessel reconstruction (Section 3.4).

Intuitively, the overall pipeline of the reconstruction process (Fig. 2) can be summarized as follows. The vessel tracking process begins with the localization of the optic disc (Section 3.6) and the segmentation of arteries and veins (Section 3.5).

Since the optic disc is the origin of all retinal vessels, its margin is used to identify the starting points of the segmented arteries and veins (Section 3.8.1). In practice, pixels at the intersection of the optic disc and the segmented vessels serve as seed points for the reconstruction process. The complete vascular structure is first reconstructed by tracing the primary vessels emerging directly from the optic disc using a growing approach (Section 3.8.2). When a branch point is encountered, an RNN model predicts the vessel's continuation (Section 3.7), ensuring the correct branch is selected for further growth. Any unselected branch points are stored as additional seed points, allowing the method to iteratively map secondary, tertiary, and higher-order vessel branches throughout the retinal vascular network. Finally, various measurements are extracted from the reconstructed vessels for characterization (Section 3.8.3).

3.1. Datasets

High-resolution fundus (HRF). The HRF dataset [56] from the Tomas Kubena Ophthalmology Clinic (Zlin, Czech Republic) provides high-resolution (3504 × 2336 pixels) fundus images. Each image is pixel-wise annotated for segmentation into four classes: background, arteries, veins, and uncertain regions. Furthermore, for each image in the dataset, there is also the annotation of the position of the optic disc. The dataset includes 45 images divided into three categories (15 images each): healthy individuals, patients with diabetic retinopathy, and patients with glaucoma. Following the official dataset split, we divide each category into training and test sets. The first five images serve as test set, while the remaining ten are used for training. In this work, to optimize the hyperparameter, we further create a validation set by removing the last two images of each category from the training data.

Retinal images vessel tree extraction (RITE). The RITE dataset [57] builds on the DRIVE dataset [58] by incorporating additional ground truth segmentation maps to distinguish between arteries and veins. Also known as AV-DRIVE or DRIVE-AV, RITE maintains the same 40 retinal images (20 training, 20 testing) as DRIVE. These images are from patients with healthy retina (33) or mild diabetic retinopathy (7). In our work, we created a validation set (4 images) from the training data for hyperparameter tuning. All images are centered on the macula and have a resolution of 768 × 584 pixels. In particular, RITE provides both the original segmentation of vessels from DRIVE and a pixel-wise classification of these vessels into arteries, veins, junctions, and uncertain regions.

FIVES. The FIVES (Fundus Image dataset for Vessel Segmentation) [59] is the largest publicly available dataset dedicated to retinal vessel segmentation (labels are provided for binary vessel segmentation, not for artery/vein classification). It includes 800 high-resolution (2048 × 2048 pixels) color fundus images evenly distributed across four categories: diabetic retinopathy, age-related macular degeneration, glaucoma, and normal retinas (200 images per category). Each image is manually annotated at the pixel level to precisely delineate the retinal vessels. The annotation process was rigorously standardized through crowdsourcing among medical experts, both junior and senior, ensuring a high degree of consistency and accuracy. Collected from 573 subjects, the dataset also features detailed quality assessments based on three critical factors: illumination and color distortion, blur, and low contrast. Each image is scored on these dimensions, which allows researchers to evaluate algorithm performance under varied real-world conditions. Notably, the inclusion of some lower-quality images – particularly from glaucoma patients whose age-related factors can affect vessel visibility – enhances the dataset's clinical relevance by reflecting the diverse challenges encountered in practice.

3.2. SegNeXt

The SegNeXt model [15] identifies key components that enhance the performance of segmentation models and integrates them into an encoder–decoder architecture for semantic segmentation. It diverges from traditional CNNs by utilizing specially designed multiscale convolutional layers instead of standard convolutions. Spatial attention is effectively incorporated into the network encoder using simple per-element multiplication. The SegNeXt network features a convolutional decoder that employs the decomposition-based Hamburger module, which extracts global context information within the decoder. Specifically, the characteristics at different levels of the decoder are combined and processed through the Hamburger module, enriching them with multiscale contextual information. SegNeXt outperforms previous segmentation methods, including transformer-based approaches. In this work, we used the tiny SegNeXt variant to for computational efficiency.

3.3. YOLO-V7

YOLOv7 [14] is a real-time object detection network that achieves a balance between high performance and efficiency. As the name "You Only Look Once" suggests, the YOLO family of models performs classification and object detection simultaneously in a single processing step for an entire image or video frame, rather than needing multiple passes or stages. YOLO models include three main components, the backbone, which extracts features from the input image, the neck, which generates feature pyramids from the extracted features, and the head, which

performs the final detection and returns the results. YOLOv7 builds on its predecessors by introducing numerous architectural changes. Compound Scaling allows for efficient scaling of the network. EE-LAN (Extended Efficient Layer Aggregation Network) improves feature extraction. A "bag of freebies" incorporates techniques that improve performance without increasing complexity. Finally, YOLOv7 employs customized loss functions, each specifically optimized for different stages of the detection process.

3.4. LSTM and BiLSTM

LSTM (Long Short Term Memory) [16] is a powerful type of artificial neural network designed specifically to overcome the limitations of traditional RNNs. Standard RNNs have difficulty in learning long-term dependencies within data sequences. This happens because information fades away as it is processed further through the network. LSTM (Long Short-Term Memory) networks manage long-term dependencies by using a gating mechanism that regulates the flow of information, allowing the network to maintain and update information over long sequences. The LSTM cell state is a vector that gets updated by the gating mechanisms. In particular, LSTMs have three special gates: the input gate, the forget gate, and the output gate. The input gate decides what new information is added to the cell state. The forget gate selects which existing information to delete from the cell. The output gate determines what information from the cell is transmitted to the rest of the network. With this control mechanism, LSTMs can selectively remember or discard information as it passes through the network. An improved version of the LSTM is the BiLSTM (Bidirectional LSTM) [17], which is essentially two LSTM networks working together. One network processes the sequence forward, while the other processes it backward. This allows BiLSTMs to capture dependencies not only within the flow of the sequence but also by considering the context from the end. The outputs of both networks are then combined to produce the final result.

3.5. Arteries and veins segmentation

Our approach uses the SegNeXt network (detailed in Section 3.2) to segment arteries and veins. The network is trained on RITE and HRF training set (Section 3.1) to distinguish between background, arteries, and veins. Hyperparameter selection is performed using the validation sets of both datasets, resulting in an optimized model that provides high-quality segmentation of arteries and veins. This segmentation is a crucial step for accurate vessel reconstruction and characterization. Section 4.1 presents the hyperparameter tuning process and the segmentation results obtained on the HRF and RITE datasets. Additionally, we assess the models' generalization by applying the segmentation networks – trained on HRF and RITE – to the FIVES dataset, which includes disease cases such as diabetic retinopathy and glaucoma. To improve generalization on FIVES, we also investigate the potential of creating an ensemble by combining models trained with diverse hyperparameter configurations. Since FIVES does not provide separate labels for arteries and veins, our evaluation focuses on quantitative metrics for overall blood vessel segmentation, while artery-vein recognition is analyzed qualitatively.

3.6. Optic disc localization

Optic disc localization employs the YOLOv7 network (details in Section 3.3) trained on the HRF training set using the validation set for early stopping. Given the absence of localization supervision, this pre-trained model is utilized for optic disc localization on the RITE and on the FIVES datasets (experiments in Section 4.2). Being able to correctly retrieve the position of the optic disc is crucial for the subsequent steps, as it serves as the starting point for vessel reconstruction.

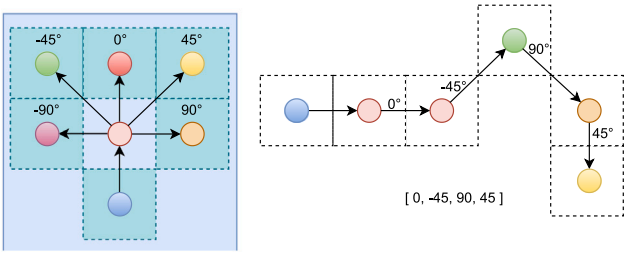

Fig. 3. An example of a vessel segment and its conversion in a sequence of directions (on the right) and the possible directions that can be traversed by the vessel (on the left).

3.7. Vessel predictor

Two RNN models, specifically LSTM and BiLSTM, are employed and compared for predicting vessel trajectories. To simplify the prediction task, we represent vessel segments as sequences of directions rather than analyzing their two-dimensional (2D) coordinates. There are two main reasons behind this approach. First, retinal images can vary significantly in resolution across devices and datasets. Using the entire 2D coordinate space as input and output to the network would be computationally expensive and require adaptation to these varying dimensions. Secondly, we can take advantage of the sequential nature of the vessels: similar to analyzing a moving piece on a chessboard, each subsequent point in a vessel segment defines a direction relative to the previous point. In practice, the vessel trajectory is determined by predicting which adjacent pixel the vessel will follow next, which defines the direction. Although there are eight possible neighboring pixels (using 8-connectivity), we chose to restrict the prediction to five directions: $-90°$, $-45°$, $0°$(straight), $45°$, and $90°$. This decision is based on two key considerations:

- Physical Plausibility — In the datasets we used (RITE, HRF, and FIVES), vessels never exhibited trajectories that reversed direction by $180°$or deviated by $\pm 135°$.
- Avoiding Improbable Predictions — Limiting the prediction to these five directions helps us avoid, a priori, the prediction of directions that are highly unlikely in practice.

Fig. 3 illustrates this conversion process for a vessel's segment.

Once the vessel segment is converted in a direction sequence, with each element one-hot encoded, the RNN predicts the sequence continuation. In particular, the RNN takes in input a sequence of directions and predicts the next element, which corresponds with one of the possible directions. The predicted output is iteratively given in input to the RNN, for a fixed number of iterations, to obtain the prediction of the next points of the vessel. To avoid inference on excessively long sequences, we set a maximum length for the input sequence. If a sequence exceeds this length, only the last segment, with the defined maximum length, is used for prediction. The vessel continuation prediction phase is summarized in Fig. 4. Finally, the predicted sequence of directions is decoded back into the vessel's trajectory in 2D coordinates, starting from the last point of the vessel.

3.8. Vessels reconstruction and characterization

The vessel reconstruction algorithm (Fig. 2) aims at reconstructing the vessel path from a starting point (seed). In our approach, the seeds are identified by finding the intersection between the vessel segmentation and the optic disc (Section 3.8.1). From each seed the reconstruction algorithm starts to "grow" a vessel, reconstructing its main path. The growth (described in 3.8.2) from the seed continues until it reaches the end of the vessel. By incorporating, as additional seed

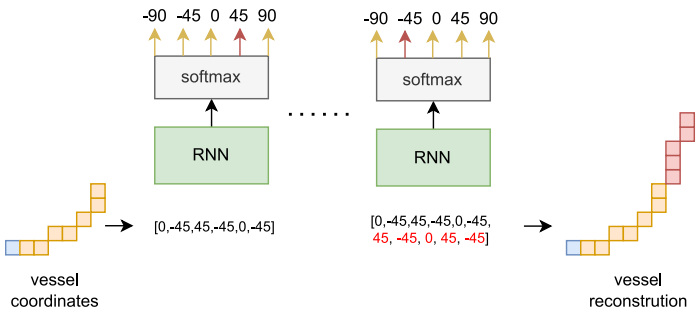

Fig. 4. Pipeline of the vessel continuation prediction: the vessel coordinates are converted in a sequence of directions used by the RNN to predict the directions of the next points that are finally reconverted to coordinates.

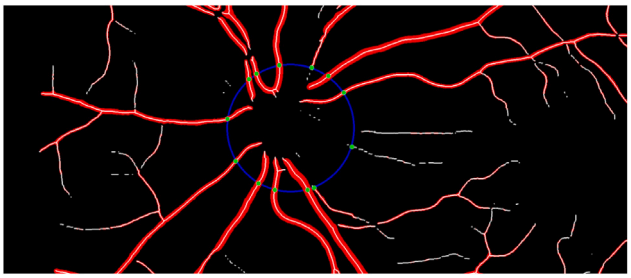

Fig. 5. An example of how the seed points are determined: the optic disc (in blue) intersects with the vessel's skeleton (in white) to find the seed points (in green).

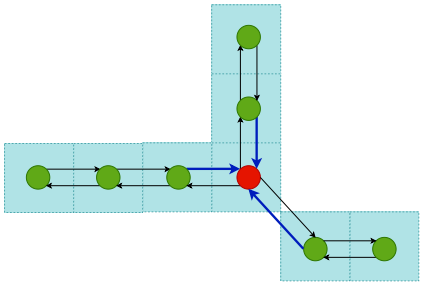

Fig. 6. Example of the procedure used to localize corners in the skeleton. The skeleton is converted into a graph whose vertices (in green) are constituted by the coordinates of the pixels (in light blue). An edge (black arrow) is added from one node to another if the pixels are close (in 8 connectivity). A vertex is an intersection (in red) if it has more than two incoming edges (blue arrows).

points, branch points not selected during the initial growing procedure, the method could iteratively map secondary, tertiary, and higher-order vessel branches across the retinal vascular network. Finally, each reconstructed vessel is characterized by extracting different quantitative measures, like its tortuosity and estimated diameter (Section 3.8.3).

3.8.1. Seeds localization

In a retinal fundus image, the optic disc, also known as the optic nerve head, is the structure where the optic nerve exits the back of the eye. It serves as the connection point between the eye and the brain, where blood vessels and nerve fibers converge as they travel to and from the light-sensitive retina. The optic disc, where blood vessels converge, serves as an ideal starting point for the reconstruction algorithm.

We determine the seeds for the initial vessel growth algorithm by finding the intersection between the optic disc and the blood vessels. Specifically, the semantic segmentation of the vessels (veins or arteries) is skeletonized, and the points where the skeleton intersects the border of the optic disc are considered seeds (see Fig. 5).

During the initial vessel reconstructions, additional seed points will be added, including the starting points of branches not selected for continuation. These seed points will then be processed to iteratively reconstruct secondary, tertiary, and higher-order vessels.

3.8.2. Vessel growth

Starting from each seed, the vessel growth algorithm is used to reconstruct the path of all the vessels. In the first phase, the algorithm builds a graph from the pixels of the vessels' skeleton: the coordinates of each pixel of the skeleton constitute a vertex, and if a pixel is close to another (in 8-connectivity), an edge between two vertices is added. The degree of the vertices in the corresponding graph is used to determine the position of the vessel's corners. More specifically, a vertex is considered a corner if it has a degree (number of incoming edges) greater than or equal to 3 (Fig. 6).

Once the corners are located, the algorithm starts to expand each vessel from the seed. The expansion process uses a region-growing algorithm that starts from the seed and continues until it reaches a corner or the end of the vessel. When the algorithm encounters a corner in the vessel's skeleton, it considers each connected branch as a potential continuation path. To identify the correct continuation branch, we propose to use a RNN network (Section 3.7). The RNN will predict the next points along the vessels' current trajectory. The branch that aligns closest with this prediction and has a similar diameter is chosen as the continuation branch. The growth proceeds from the first point of the selected branch, while the starting points of the branches not chosen during the vessel reconstruction are added to the seed list. These additional seed points can then be used to reconstruct secondary, tertiary, and higher-order vessels. The procedure is repeated until no valid expansion points are found. A pseudocode of the algorithm is given in Algorithm 1.

Algorithm 1 Vessel Growth Algorithm

```
seeds ← GetIntersections(optic disc, skeleton)
while |seeds| > 0 do
    seed ← Pop(seeds)        ▷ Take a seed from seeds and remove it
    vessel ← ∅
    starting_point ← seed
    repeat
        vessel ← vessel ∪ RegionGrowing(starting_point)
        candidates ← SearchExpansionCandidates(vessel)
        if not IsEmpty(candidates) then
            starting_point ← ClosestCandidate(candidates)
            seeds ← seeds ∪ OtherCandidates(candidates)
        end if
    until IsEmpty(candidates)
end while
```

Segmentation of vessels, particularly for those that are very small or close to branch points, can often be discontinuous. These unavoidable

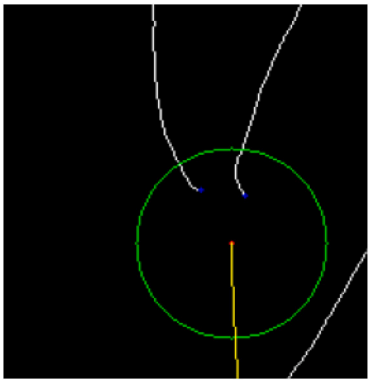

Fig. 7. Vessel growth with "broken vessels": From the last point (in red) of the incoming vessel (in yellow), a circular region with a fixed radius (in green) is considered. The closest points of the other skeleton connected components (in blue) are considered as new possible expansion points.

gaps can mislead the vessel growth process and cause it to stop prematurely. To solve this problem and improve robustness, our algorithm implements the following strategy. If no valid expansion point is found, the algorithm searches a circular neighborhood (with a predefined radius of 30 pixels[1]) around the vessel's endpoint. If points belonging to different skeleton branches are located within this neighborhood (see Fig. 7), they are treated as potential continuation, similar to the branches connected to a corner. Using these points, the vessel growth procedure resumes as before. Otherwise, if no suitable points are found, the reconstruction of the current vessel is considered complete.

3.8.3. Vessel characterization

This section describes the approach used to characterize the reconstructed vessels. The proposed vessel reconstruction method allows for comprehensive characterization, encompassing the entire vessel path rather than just isolated segments. This analysis would be impossible without the reconstruction obtained by the presented approach. The reconstruction process returns the vessels as lists of consecutive 2D coordinates. The first point corresponds to the initial seed at the intersection with the optic disc, and the last point marks the opposite end point of the vessel. Disconnected segments appear during reconstruction between consecutive branches that are merged together, resulting in gaps within the vessel. To solve this problem, the characterization phase fills these gaps using a shortest path algorithm, effectively connecting the missing sections with a set of 2D coordinates. As a result of the reconstruction phase, all the vessel's points are ordered based on their distance from the starting seed. This sorting facilitates the extraction of various vessel metrics for further analysis. The main characteristics obtained from the vessels include the length, the diameter, and some tortuosity measures. To compute the tortuosity, we employ the common metrics defined in [60]: the Distance Metric (DM), the Summation of Angles Metric (SOAM), and the Inflection Count Metric (ICM).

Distance metric (DM). The ratio of the arc over the chord. It is the most common measure used in the scientific literature.

$$DM = \frac{Arc}{Chord} \tag{1}$$

where Arc is the actual path length of the curve and $Chord$ is the linear distance between the endpoints.

Summation of angles metric (SOAM). The summation of the angles (α_i) between two vectors created by three points of the skeletonized segment, normalized by the Arc.

$$SOAM = \sum_{i=1}^{n} \frac{180 - \alpha_i}{Arc} \tag{2}$$

Inflection count metric (ICM). ICM is calculated by counting the points where the curve changes concavity (inflection points) and normalized by multiplying it by DM.

$$ICM = (iPoints + 1) \cdot DM \tag{3}$$

where $iPoints$ is the total number of inflection points and DM is the distance metric (Eq. (1)).

4. Experiments and results

This Section presents the experimental setup and the results of the proposed approach. In particular, Sections 4.1 and 4.2 describe the experimental setups and the results of veins/arteries segmentation and optic disc localization. Instead, in Sections 4.3 and 4.4 we provide an overview of the experiments conducted for vessel prediction and characterization.

4.1. Segmentation network

This section details the experimental setup and results for vein and artery segmentation. We trained the tiny SegNeXt model using the publicly available implementation from GitHub.[2] Experiments were conducted on both HRF and RITE datasets. To identify the optimal training configuration, we systematically evaluate the impact of individual hyperparameters on model performance. We employed a learning rate of 10^{-6}, a batch size of 2 with gradient accumulation every 8 iterations and we conducted a series of experiments in which we changed a single hyperparameter at a time.

In particular, we evaluate the impact of:

- Loss functions — cross entropy (EC) and its combination with the dice score (EC + dice)
- Image flip — Horizontal (H), Vertical (V), Horizontal + Vertical (H+V) and None
- Rotation — [90°, 180°and 270°], random angle in the range [+15°, −15°], and None
- Crop size — we explored different values, for HRF (1024, 2048) and for RITE (512, 1024) due to different original image sizes. To analyze the entire image we employed a sliding window approach with windows that overlap by 1/4 of the crop size.
- Multi-scale factor — we tested factors 1, 1.2 and 1.4 on HRF and 1, 4 and 5 on RITE to incorporate data augmentation at different scales.

Table 2 reports the results obtained using different hyperparameters on the validation set of HRF and RITE. We employed the mean Intersection over Union (mIoU) to evaluate the segmentation performance in recognizing arteries, veins, and background.

We evaluated the top-performing models for each dataset on the test set. The best model achieved an average mIoU of 75.51% on HRF and 75.09% on RITE for artery/vein segmentation. To further improve performance, we also investigated an ensemble approach that combines, using a max-value fusion approach, the predictions of the top three models listed in Table 2 for HRF and RITE. This ensemble method led to a notable increase in segmentation accuracy, achieving an average mIoU of 79.60% on HRF and 79.66% on RITE. In Fig. 8 a qualitative evaluation of the segmentation results obtained on HRF and RITE is reported.

[1] This hyperparameter was selected empirically based on visual evaluation.

[2] https://github.com/visual-attention-network/segnext

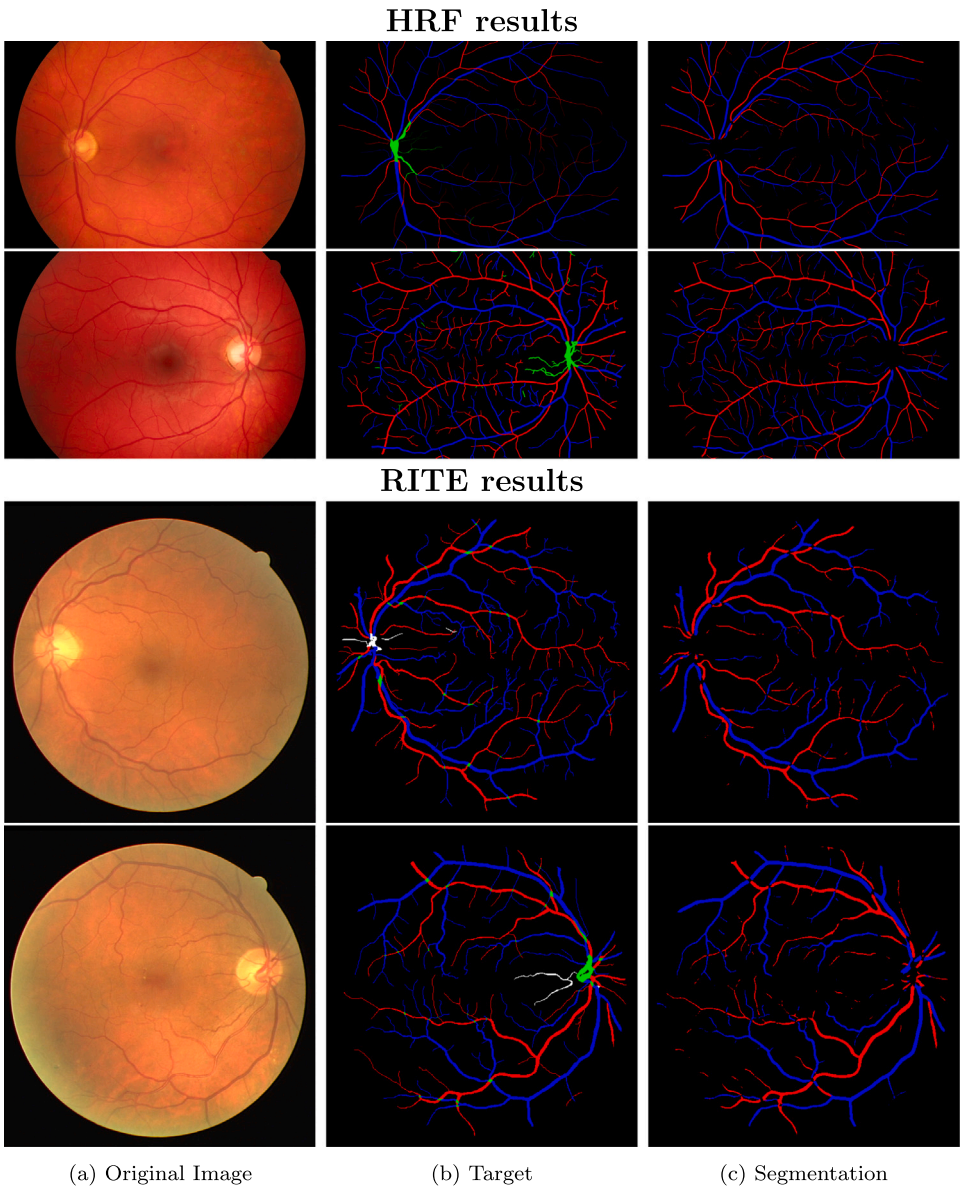

Fig. 8. Segmentation results on some images from HRF (top) and RITE (bottom), in red the arteries and in blue the veins.

To compare our performance with existing methods, we used one-vs-one evaluation metrics (sensitivity, specificity, and accuracy) for both the artery/vein (A/V) and blood vessel/fundus (BV/BG) classification tasks. Arteries and blood vessels were considered the positive class for each task, respectively. These metrics have been commonly used in the literature to evaluate performance on HRF and RITE datasets. In Table 3, we compare our results with other relevant approaches from the literature. Our model achieves comparable results with the state-of-the art on A/V classification on the HRF dataset. For BV/BG classification on HRF, our results are very competitive with the current leading methods. While our model does not achieve the best results on the RITE dataset, it still provides excellent segmentation quality, approaching state-of-the-art performance.

To further evaluate the model's generalization ability, we test the top-performing models trained on HRF and RITE on the test set of the FIVES dataset. Additionally, we investigate ensemble methods to improve segmentation accuracy. Specifically, we construct two ensemble models by applying a max-value fusion approach to the output probability maps of the three best models trained on HRF and the three best models trained on RITE, respectively. Since FIVES lacks ground truth for artery-vein segmentation, our quantitative analysis focuses on overall blood vessel segmentation (Table 4).

The results show that the ensemble of the three models trained on RITE achieves strong segmentation performance on the FIVES dataset, which features retinal images with previously unseen pathologies. Finally, Fig. 9 presents a qualitative evaluation of artery and vein segmentation achieved using this ensemble. Images representing three critical retinal conditions are presented: age-related macular degeneration, diabetic retinopathy, and glaucoma. The strong results highlight the ensemble's ability to adapt and effectively analyze previously unseen pathologies.

4.2. Localization network

For optic disc localization, we employed the YOLOv7 network (using the implementation available from GitHub[3]), trained on the HRF dataset, which provides optic disc annotations for each image (center

[3] https://github.com/WongKinYiu/yolov7

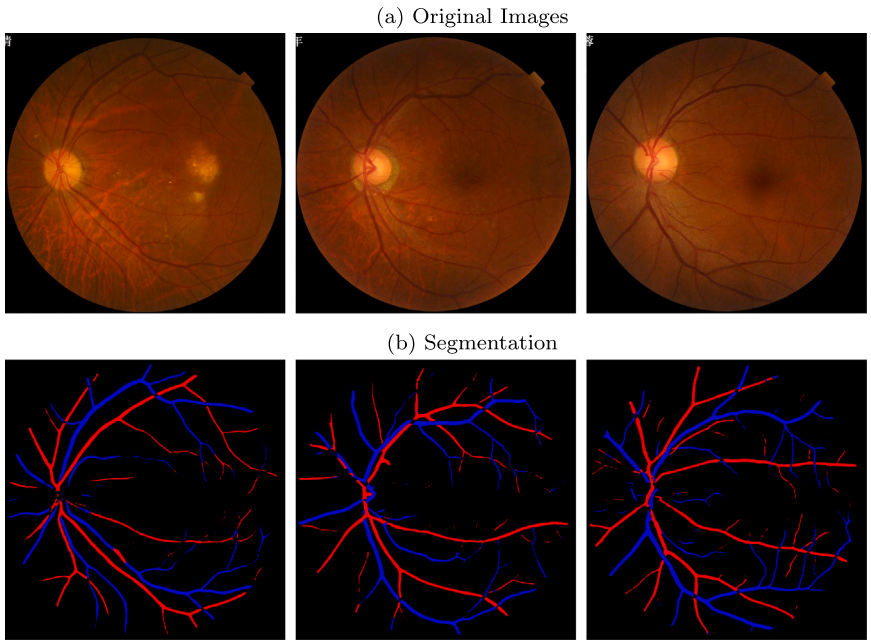

Fig. 9. Results from the ensemble of the three best models trained on RITE for artery and vein segmentation on the FIVES test set. In the images, we can observe different retinal pathologies: age-related macular degeneration (left), diabetic retinopathy (center), and glaucoma (right).

Table 2
Comparison of the models trained using different hyperparameters on the validation set of HRF and RITE.

HRF					
Loss	Scale	Crop	Rotation	Flip	Results
CE	1	1024	90, 180, 270	H	76.69%
CE+Dice	1	1024	90, 180, 270	H	76.80%
CE+Dice	**1.2**	**1024**	**90, 180, 270**	H	**76.90%**
CE+Dice	1.4	1024	90, 180, 270	H	76.88%
CE+Dice	1.2	2048	90, 180, 270	H	76.84%
CE+Dice	1.2	1024	+15,−15	H	76.08%
CE+Dice	1.2	1024	None	H	75.91%
CE+Dice	1.2	1024	90, 180, 270	H+V	76.73%
CE+Dice	1.2	1024	90, 180, 270	None	76.70%
CE+Dice	1.2	1024	90, 180, 270	V	76.86%
RITE					
Loss	Scale	Crop	Rotation	Flip	Results
CE	1	512	90, 180, 270	H	65.83%
CE+Dice	1	512	90, 180, 270	H	67.17%
CE+Dice	**4**	**512**	**90, 180, 270**	H	**74.56%**
CE+Dice	5	512	90, 180, 270	H	74.15%
CE+Dice	4	1024	90, 180, 270	H	73.88%
CE+Dice	4	512	+15, −15	H	73.94%
CE+Dice	4	512	None	H	73.36%
CE+Dice	4	512	90, 180, 270	H+V	74.19%
CE+Dice	4	512	90, 180, 270	None	74.45%
CE+Dice	4	512	90, 180, 270	V	74.50%

point and radius). The default hyperparameters were used, and early stopping was applied during training based on the validation set. The performance of the YOLOv7 model was quantitatively evaluated on the HRF test set. The model achieved a very high mean Average Precision (mAP@0.5) of 0.934, underscoring the model's accuracy in detecting the optic disc. Instead, since the RITE and the FIVES datasets lack optic disc annotations, we employed the model trained on the HRF dataset for optic disc localization on RITE and FIVES images. Fig. 10 provides a qualitative assessment of the localization performance on HRF, RITE, and FIVES datasets. Notably, the system demonstrates strong generalization capabilities, effectively adapting to FIVES images containing various pathological conditions.

4.3. Vessel reconstruction

In the vessel reconstruction phase, a RNN is employed to predict the continuation of vascular segments and reconstruct the primary vascular network within the fundus images. To train and validate the RNN, a dataset of vessel segments is extracted from the HRF and RITE datasets using the segmentation masks. Given the absence of explicit supervision for the main vessel paths in each image, we simplify the problem by segmenting the overall vessel tree into smaller segments. These segments are defined by the connected components of the skeleton that are separated by corners. The segment coordinates are transformed in a sequence of directions, as described in Section 3.8. Through this process, a dataset of sequence directions is generated (3068 sequences for HRF and 2026 for RITE), representing individual vessel segments. This dataset is divided in training, validation, and test sets, adhering to the original split of HRF and RITE. The obtained datasets are then used to train and evaluate the vessels's prediction module, separately for the HRF and RITE. The RNN is specifically designed to predict the next element in a sequence based on previous directions. In this work, we evaluate two RNN architectures, LSTM and BiLSTM, using a grid search approach to test different combinations of hyperparameters for each model. Furthermore, we studied the influence of sequence length on model performance by setting a maximum length for each sequence and analyzing how this constraint affects the results. Table 5 summarizes the results, including the optimal hyperparameters identified via grid search for both datasets and the influence of sequence length.

The results indicate that a sequence length of 30 consistently yields the best performance across all configurations. Additionally, in the best configuration, the BiLSTM provides a slight improvement in performance. The best model for both datasets is also evaluated on the test set as a regression task. We converted the predicted directions back into pixel space and compared them to the vessel's actual next point using the Mean Square Error (MSE), Mean Absolute Error (MAE), and Root Mean Square Error (RMSE) (see Table 6). Additionally, we reported accuracy for completeness; however, given the complexity of vessel continuation – where multiple plausible paths can exist – MSE offers a more informative evaluation by capturing how closely the predicted continuation aligns with the actual vessel trajectory.

The best model, as described in Section 3.8, is used to reconstruct the vessel paths. At the end of the reconstruction phase, we obtain the

Table 3
Comparison with the state-of-the-art results on the test set of HRF and RITE dataset.

HRF						
Method	A/V Classification			BV/BG Classification		
	Sens.	Spec.	Acc.	Sens.	Spec.	Acc.
Chen et al. [43]	97.0%	97.2%	97.1%	78.1%	98.2%	96.5%
Galdran et al. [44]	98.1%	93.1%	95.3%	81.1%	98.1%	96.7%
Karlsson and Hardarson [45]	97.0%	96.5%	96.7%	**86.1%**	97.0%	96.1%
Hu et al. [46]	93.3%	95.3%	94.4%	69.0%	**99.0%**	96.2%
Morano et al. [10]	97.9%	97.7%	97.8%	82.7%	97.8%	96.6%
Mokan et al. [47]	89.0%	89.0%	89.0%	–	–	–
Wan et al. [48]	99.1%	99.3%	99.26%	–	–	–
Chen et al. [49]	**99.3%**	**99.4%**	**99.3%**	–	–	–
Our	98.8%	98.5%	98.7%	78.4%	98.6%	97.1%
Our Ensamble best 3 models	98.9%	98.7%	98.8%	78.9%	98.7%	**97.2%**
RITE						
Method	A/V Classification			BV/BG Classification		
	Sens.	Spec.	Acc.	Sens.	Spec.	Acc.
Girard et al. [39]	86.3%	86.6%	86.5%	78.4%	98.1%	95.7%
Galdran et al. [40]	89.0%	90.0%	89%	**94.0%**	93.0%	93.0%
Ma et al. [38]	93.4%	95.5%	94.5%	79.1%	98.1%	95.7%
Hemelings et al. [37]	95.1%	92.7%	93.8%	77.6%	**98.7%**	96.0%
Morano et al. [42]	87.4%	90.8%	89.2%	79.1%	98.6%	96.1%
Galdran et al. [44]	88.8%	96.0%	92.7%	83.0%	98.1%	96.2%
Karlsson and Hardarson [45]	95.1%	96.0%	95.6%	82.2%	97.6%	95.6%
Chen et al. [43]	95.3%	97.2%	96.3%	81.5%	97.8%	95.7%
Hu et al. [46]	93.3%	95.3%	94.4%	79.0%	98.1%	95.6%
Morano et al. [10]	**95.7%**	97.3%	**96.6%**	80.1%	98.6%	96.2%
Mokan et al. [47]	91.0%	87.1%	89.0%	–	–	–
Wan et al. [48]	94.1%	94.2%	94.1%	–	–	–
Chen et al. [49]	95.4%	**97.5%**	96.5%	–	–	–
Our	93.7%	95.6%	94.7%	81.2%	98.6%	97.2%
Our Ensamble best 3 models	94.4%	95.1%	94.7%	81.8%	**98.7%**	**97.3%**

Table 4
Blood vessel segmentation results of the models trained on HRF and RITE evaluated on the FIVES test set.

Model	mIoU	Sens.	Spec.	Acc.
Best model on HRF	73.07%	51.75%	99.69%	96.50%
Best model on RITE	75.91%	61.70%	99.08%	96.75%
Ensemble best 3 models on HRF	74.00%	52.89%	99.79%	96.68%
Ensemble best 3 models on RITE	**76.17%**	**61.33%**	**99.20%**	**96.87%**

Table 5
Results on the HRF and RITE validation set for the vessel continuation prediction network.

HRF						
Len.	Batch	Lr.	Drop.	Units	Model	Acc.
20	256	0.001	0.0	16	LSTM	72.86%
30	512	0.001	0.0	32	LSTM	73.08%
40	256	0.0001	0.0	64	LSTM	68.59%
20	512	0.01	0.2	64	BiLSTM	68.08%
30	**512**	**0.001**	**0.2**	**256**	**BiLSTM**	**76.19%**
40	1024	0.01	0.2	128	BiLSTM	68.61%
RITE						
Len.	Batch	Lr.	Drop.	Units	Model	Acc.
20	256	0.001	0.0	16	LSTM	72.86%
30	512	0.001	0.0	32	LSTM	73.08%
40	512	0.01	0.0	8	LSTM	72.87%
20	1024	0.01	0.0	128	BiLSTM	72.79%
30	**512**	**0.01**	**0.0**	**16**	**BiLSTM**	**73.10%**
40	256	0.01	0.1	32	BiLSTM	72.60%

paths of all the vessels (arteries and veins) in the image. An example of the reconstructed vessels is provided in Fig. 11. Notably, the proposed method demonstrates strong generalization capabilities, even on FIVES images – representing various pathologies – which were not included in the training set.

4.4. Vessel characterization

The vessel characterization is performed on the reconstructed vessels, extracting features like the vessel diameter, the length, and the tortuosity. Since ground truth supervision for the tortuosity estimation is not present, we report some qualitative results in Fig. 12.

Although the vessels analyzed have similar lengths, they show various degrees of tortuosity using different metrics. Vessel (c) appears to be the least tortuous, as evidenced by all the calculated tortuosity indices. Instead, vessel (a) has a higher SOAM, indicating larger angular changes than (b). In contrast, vessel (b) has a higher ICM, suggesting more frequent local changes in the concavity of the curve. Interestingly, the SOAM and ICM indices provide additional information on the characteristics of vessels (a) and (b). This demonstrates that, thanks to the proposed approach, it is possible to extract significant measurements capable of providing a comprehensive understanding of the vessel's characteristics.

5. Discussion

In this work, we introduce a comprehensive tool for retinal image vessel segmentation, reconstruction, and characterization. By leveraging an enhanced, automatic reconstruction of retinal vasculature, our approach may facilitate a deeper understanding of the underlying disease mechanisms, thereby enabling more targeted treatment strategies. Its integration into clinical workflows could not only alleviate the workload on healthcare professionals but also support large-scale screening and rapid patient assessments, optimizing resource allocation and improving overall patient care. Moreover, the automatic extraction of retinal vascular patterns could provide a reliable method for tracking morphological changes over time — changes that are closely associated with a spectrum of ocular diseases, including diabetic retinopathy, age-related macular degeneration, and glaucoma [61,62]. The experimental results demonstrate the effectiveness of our comprehensive approach to retinal vessel analysis. While each individual component of our

Table 6
Results of the best models on the HRF and RITE test set for the vessel continuation prediction network.

HRF				
Model	Accuracy	MSE	MAE	RMSE
BiLSTM	72.15%	1.78×10^{-6}	1.76×10^{-6}	4.73×10^{-6}
RITE				
Model	Accuracy	MSE	MAE	RMSE
BiLSTM	69.51%	4.81×10^{-5}	4.70×10^{-5}	1.20×10^{-4}

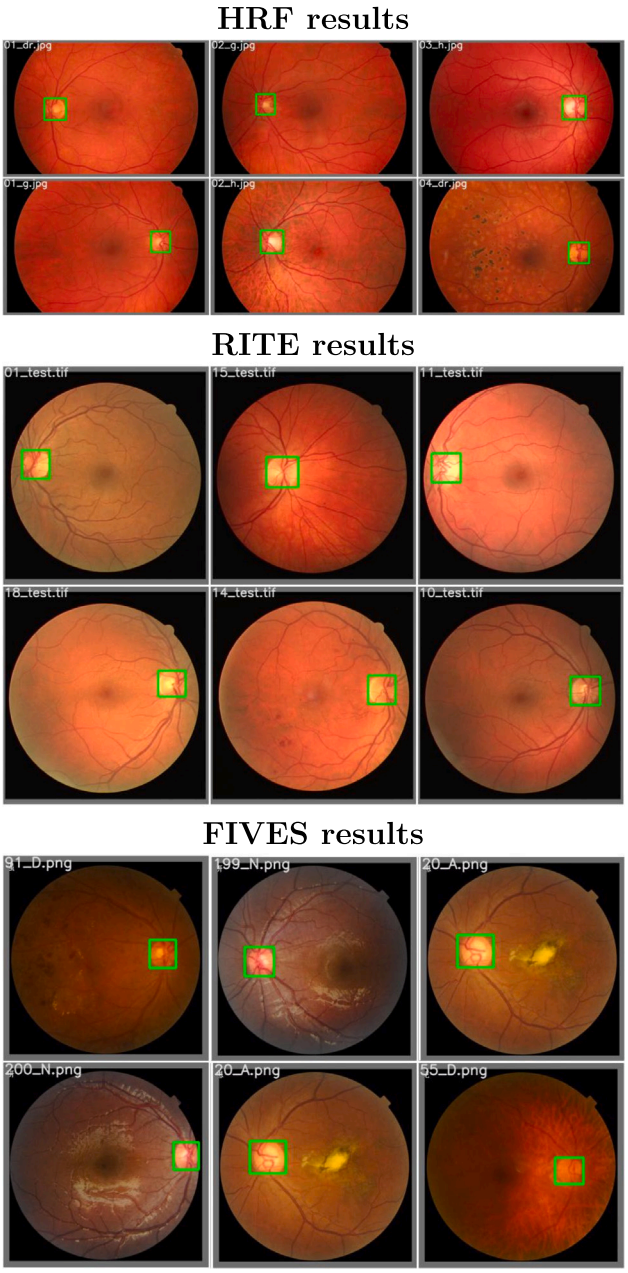

Fig. 10. Example results of optic disc localization.

system – including vessel segmentation, optic disc localization, and vessel reconstruction – achieves performance comparable to state-of-the-art methods, the significant contribution of this work lies in the successful integration of these components into a unified pipeline for complete fundus image analysis. Our segmentation network shows strong performance in discriminating between arteries and veins while accurately detecting vessel structures. It also exhibits strong generalization ability on a heterogeneous dataset of unseen retinal pathology images. The optic disc localization component demonstrates robust detection capabilities, even generalizing well to datasets without explicit annotations. The vessel reconstruction module effectively predicts vessel continuations, enabling the tracing of primary vascular pathways throughout the fundus image. However, the true value of our approach emerges from the synergy between these components. By combining artery/vein segmentation with vessel path reconstruction, we enable the extraction of quantitative measurements that would be impossible to obtain from any single component alone. The integration allows for comprehensive vessel characterization, including precise measurements of vessel diameter, length, and various tortuosity metrics. These metrics provide complementary information about vessel morphology, offering a more complete picture of vascular health. The potential clinical impact of this work is significant. By providing objective, quantitative measurements of retinal vessels, our system could help reduce the subjectivity inherent in human interpretation of fundus images. This quantification enables precise tracking of changes in retinal structures over time, potentially allowing for earlier detection of pathological changes. Furthermore, the automated analysis could provide valuable decision support for clinicians, potentially speeding up diagnosis times and enabling more precise and early detection of various retinal conditions. The vessel characterization capabilities demonstrated in our results suggest promising applications in clinical practice. Different tortuosity metrics capture distinct aspects of vessel morphology, providing clinicians with detailed information about vascular health. It is important to note that once each vessel is reconstructed – meaning all its pixels and junction points are identified – it becomes possible to extract a comprehensive set of measurements and geometrical features. While this paper presents a selection of these metrics, the approach is not limited to them and can be extended to additional analyses as needed. This comprehensive characterization could be particularly valuable for monitoring disease progression and treatment response.

The computational efficiency of the proposed approach primarily depends on the hardware used and the resolution of the images. The segmentation of arteries and veins, as well as optic disc localization, operates almost in real-time (under 1 s per image) on an NVIDIA GeForce RTX 4090. However, the vessel characterization process is significantly more computationally demanding. Its execution time varies based on the number of vessels, averaging 26 s for RITE images and 89 s for HRF images. Although this process is not real-time, it does not compromise the system's usefulness in a clinical setting. The ability to provide detailed vascular characterization within a minute remains well within an acceptable timeframe for diagnostic workflows. Given that such analyses are typically performed offline or as part of pre-consultation assessments, the system can seamlessly integrate into clinical routines, aiding in early disease detection, monitoring, and treatment planning without causing workflow disruptions.

However, some limitations should be acknowledged. While our system shows promising results on standard datasets, it has not yet been validated in a real clinical setting. Future work should include clinical trials in collaboration with ophthalmologists to verify the system's effectiveness in practical usage scenarios. Such trials would be crucial for understanding how the system performs across different patient populations and clinical conditions, and how it can best be

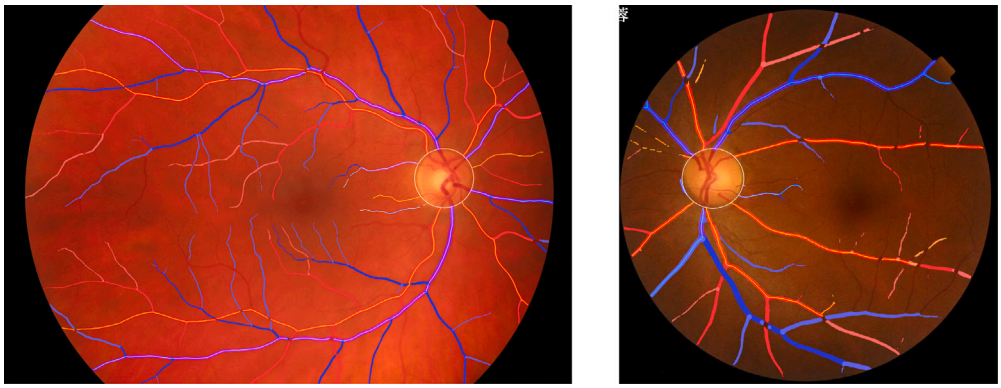

(a) HRF (b) FIVES

Fig. 11. Reconstructed vessels. Arteries are depicted in varying shades of red, while veins are rendered in different blue tones. The tones indicate different levels of the vessel structure. The images come from two datasets: (a) an image from HRF and (b) an image from FIVES, with the FIVES image derived from a glaucoma case.

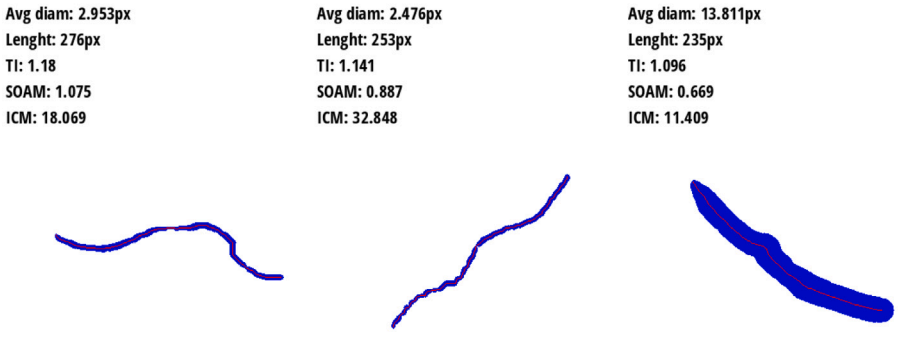

(a) (b) (c)

Fig. 12. Examples of vessel characterization. At the bottom of each image, the analyzed reconstruction of the vessel segmentation (in blue) and the corresponding center lines (in red) are shown. On top, the average diameter (Avg diam), the vessel length (length), and the tortuosity measures (TI, SOAM, ICM).

integrated into clinical workflows. Additionally, it would be beneficial to validate the vessel characterization metrics with the guidance of medical experts. Another limitation is that the accuracy of the reconstruction, and consequently the extracted metrics, is highly dependent on both the segmentation quality and the vessel continuation predictor. Improving the performance of these two modules will be a focus of future work. However, it is important to note that the segmentation algorithm operates independently of the reconstruction system, allowing for seamless integration with any existing network trained for artery and vein segmentation with higher precision. Looking forward, several directions for future research emerge. The system could be enhanced with additional biomarkers and characteristics relevant to specific pathologies. Integration with longitudinal patient data could enable predictive modeling of disease progression. Finally, future development will focus on designing a user interface that enhances interpretability and supports clinical practice. This advancement could significantly improve usability by helping clinicians more effectively interpret and utilize the quantitative measurements provided by the system.

6. Conclusions

This paper presents a comprehensive method for characterizing retinal blood vessels, leveraging advanced deep learning models and ad hoc algorithms to go beyond simple vessel segmentation. The developed algorithm effectively reconstructs the vessels' path, allowing their individual characterization. More specifically, the process begins by locating the optic disc and segmenting the arteries and veins, using the disc's margin as the origin for vessel reconstruction. Primary vessels are traced out from the disc using a growing approach, while an RNN model predicts the continuation at branch points. Any unselected branches are stored as seed points, enabling the iterative mapping of secondary, tertiary, and higher-order vessel branches. The system organically integrates neural networks such as YOLO for optic disc localization, SegNeXt for vessel segmentation, and bidirectional LSTM for a comprehensive analysis of vascular structures in retinal images. Our system's significant contribution is its ability to extract features from the vessels, which are individually reconstructed, that would otherwise be unfeasible to analyze. The capability to measure geometrical features such as tortuosity and vessel diameter can provide significant clinical insights, potentially improving the diagnosis and monitoring of diseases manifested through vascular changes. Our approach promises to be a valuable tool in ophthalmology, particularly for the early diagnosis and treatment planning of retinal and systemic vascular conditions. Future work will focus on clinical trials in collaboration with ophthalmologists to assess the system's effectiveness in real-world scenarios and validate vessel characterization metrics with medical experts. Efforts will also be directed towards improving segmentation and vessel continuation prediction, integrating additional biomarkers, leveraging longitudinal patient data for predictive modeling, and integrating the proposed tool into a decision support system to assist ophthalmologists.

CRediT authorship contribution statement

Paolo Andreini: Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Simone Bonechi:** Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Ethics in publishing statement

This research presents an accurate account of the work performed, all data presented are accurate and methodologies detailed enough to permit others to replicate the work. This manuscript represents entirely original works and or if work and/or words of others have been used, that this has been appropriately cited or quoted and permission has been obtained where necessary. This material has not been published in whole or in part elsewhere. That generative AI and AI-assisted technologies have not been used to create or alter images. All authors have been personally and actively involved in substantive work leading to the manuscript and will hold themselves jointly and individually responsible for its content.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] M.M. Fraz, P. Remagnino, A. Hoppe, B. Uyyanonvara, A.R. Rudnicka, C.G. Owen, S.A. Barman, Blood vessel segmentation methodologies in retinal images–a survey, Comput. Methods Programs Biomed. 108 (1) (2012) 407–433.

[2] D.D. Patil, R.R. Manza, Design new algorithm for early detection of primary open angle glaucoma using retinal optic cup to disc ratio, in: 2016 International Conference on Electrical, Electronics, and Optimization Techniques, ICEEOT, IEEE, 2016, pp. 148–151.

[3] J.J. Kanski, B. Bowling, Clinical Ophthalmology: A Systematic Approach, Elsevier Health Sciences, 2011.

[4] M.D. Abràmoff, M.K. Garvin, M. Sonka, Retinal imaging and image analysis, IEEE Rev. Biomed. Eng. 3 (2010) 169–208.

[5] R. Biswas, A. Vasan, S.S. Roy, Dilated deep neural network for segmentation of retinal blood vessels in fundus images, Iran. J. Sci. Technol. Trans. Electr. Eng. 44 (1) (2020) 505–518.

[6] A. Budai, J. Odstrcilik, R. Kolar, J. Hornegger, J. Jan, T. Kubena, G. Michelson, A public database for the evaluation of fundus image segmentation algorithms, Investig. Ophthalmol. Vis. Sci. 52 (14) (2011) 1345–1345.

[7] A.E. Ilesanmi, T. Ilesanmi, A.G. Gbotoso, A systematic review of retinal fundus image segmentation and classification methods using convolutional neural networks, Heal. Anal. (2023) 100261.

[8] L. Tong, T. Li, Q. Zhang, Q. Zhang, R. Zhu, W. Du, P. Hu, LiViT-Net: A U-net-like, lightweight transformer network for retinal vessel segmentation, Comput. Struct. Biotechnol. J. 24 (2024) 213–224.

[9] P. Andreini, G. Ciano, S. Bonechi, C. Graziani, V. Lachi, A. Mecocci, A. Sodi, F. Scarselli, M. Bianchini, A two-stage GAN for high-resolution retinal image generation and segmentation, Electronics 11 (1) (2021) 60.

[10] J. Morano, G. Aresta, H. Bogunović, Rrwnet: Recursive refinement network for effective retinal artery/vein segmentation and classification, 2024, arXiv preprint arXiv:2402.03166.

[11] N. Ahmad, K.-T. Lai, M. Tanveer, Retinal blood vessel tracking and diameter estimation via Gaussian process with rider optimization algorithm, IEEE J. Biomed. Heal. Inform. (2023).

[12] H. Guedri, M. Ben Abdallah, F. Echouchene, H. Belmabrouk, Novel computerized method for measurement of retinal vessel diameters, Biomedicines 5 (2) (2017) 12.

[13] M.E. Asl, N.A. Koohbanani, A.F. Frangi, A. Gooya, Tracking and diameter estimation of retinal vessels using Gaussian process and radon transform, J. Med. Imaging 4 (3) (2017) 034006–034006.

[14] C.-Y. Wang, A. Bochkovskiy, H.-Y.M. Liao, YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 7464–7475.

[15] M.-H. Guo, C.-Z. Lu, Q. Hou, Z. Liu, M.-M. Cheng, S.-M. Hu, Segnext: Rethinking convolutional attention design for semantic segmentation, Adv. Neural Inf. Process. Syst. 35 (2022) 1140–1156.

[16] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Comput. 9 (8) (1997) 1735–1780.

[17] A. Graves, S. Fernández, J. Schmidhuber, Bidirectional LSTM networks for improved phoneme classification and recognition, in: International Conference on Artificial Neural Networks, Springer, 2005, pp. 799–804.

[18] M. Ben Abdallah, A.T. Azar, H. Guedri, J. Malek, H. Belmabrouk, Noise-estimation-based anisotropic diffusion approach for retinal blood vessel segmentation, Neural Comput. Appl. 29 (2018) 159–180.

[19] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18, Springer, 2015, pp. 234–241.

[20] Y. Fu, J. Chen, J. Li, D. Pan, X. Yue, Y. Zhu, Optic disc segmentation by U-net and probability bubble in abnormal fundus images, Pattern Recognit. 117 (2021) 107971.

[21] X. Meng, X. Xi, L. Yang, G. Zhang, Y. Yin, X. Chen, Fast and effective optic disk localization based on convolutional neural network, Neurocomputing 312 (2018) 285–295.

[22] S. Maiti, D. Maji, A.K. Dhara, G. Sarkar, Automatic detection and segmentation of optic disc using a modified convolution network, Biomed. Signal Process. Control. 76 (2022) 103633.

[23] R. Bhattacharya, R. Hussain, A. Chatterjee, D. Paul, S. Chatterjee, D. Dey, PY-net: Rethinking segmentation frameworks with dense pyramidal operations for optic disc and cup segmentation from retinal fundus images, Biomed. Signal Process. Control. 85 (2023) 104895.

[24] Y. Zhang, M. He, Z. Chen, K. Hu, X. Li, X. Gao, Bridge-net: Context-involved U-net with patch-based loss weight mapping for retinal blood vessel segmentation, Expert Syst. Appl. 195 (2022) 116526.

[25] S. Lian, L. Li, G. Lian, X. Xiao, Z. Luo, S. Li, A global and local enhanced residual u-net for accurate retinal vessel segmentation, IEEE/ ACM Trans. Comput. Biol. Bioinform. 18 (3) (2019) 852–862.

[26] X. Li, Y. Jiang, M. Li, S. Yin, Lightweight attention convolutional neural network for retinal vessel image segmentation, IEEE Trans. Ind. Inform. 17 (3) (2020) 1958–1967.

[27] D. Wang, A. Haytham, J. Pottenburgh, O. Saeedi, Y. Tao, Hard attention net for automatic retinal vessel segmentation, IEEE J. Biomed. Heal. Inform. 24 (12) (2020) 3384–3396.

[28] Y. Cheng, M. Ma, L. Zhang, C. Jin, L. Ma, Y. Zhou, Retinal blood vessel segmentation based on densely connected U-net, Math. Biosci. Eng. 17 (4) (2020) 3088–3108.

[29] H. Wu, W. Wang, J. Zhong, B. Lei, Z. Wen, J. Qin, Scs-net: A scale and context sensitive network for retinal vessel segmentation, Med. Image Anal. 70 (2021) 102025.

[30] Y. Wu, Y. Xia, Y. Song, Y. Zhang, W. Cai, NFN+: A novel network followed network for retinal vessel segmentation, Neural Netw. 126 (2020) 153–162.

[31] C. Guo, M. Szemenyei, Y. Yi, W. Wang, B. Chen, C. Fan, Sa-unet: Spatial attention u-net for retinal vessel segmentation, in: 2020 25th International Conference on Pattern Recognition, ICPR, IEEE, 2021, pp. 1236–1242.

[32] P. Costa, A. Galdran, M.I. Meyer, M. Niemeijer, M. Abràmoff, A.M. Mendonça, A. Campilho, End-to-end adversarial retinal image synthesis, IEEE Trans. Med. Imaging 37 (3) (2017) 781–791.

[33] A. Beers, J. Brown, K. Chang, J.P. Campbell, S. Ostmo, M.F. Chiang, J. Kalpathy-Cramer, High-resolution medical image synthesis using progressively grown generative adversarial networks, 2018, arXiv preprint arXiv:1805.03144.

[34] G. Lepetit-Aimon, R. Duval, F. Cheriet, Large receptive field fully convolutional network for semantic segmentation of retinal vasculature in fundus images, in: Computational Pathology and Ophthalmic Medical Image Analysis: First International Workshop, COMPAY 2018, and 5th International Workshop, OMIA 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16-20, 2018, Proceedings 5, Springer, 2018, pp. 201–209.

[35] S. Zhang, R. Zheng, Y. Luo, X. Wang, J. Mao, C.J. Roberts, M. Sun, Simultaneous arteriole and venule segmentation of dual-modal fundus images using a multi-task cascade network, IEEE Access 7 (2019) 57561–57573.

[36] H. Kang, Y. Gao, S. Guo, X. Xu, T. Li, K. Wang, Avnet: A retinal artery/vein classification network with category-attention weighted fusion, Comput. Methods Programs Biomed. 195 (2020) 105629.

[37] R. Hemelings, B. Elen, I. Stalmans, K. Van Keer, P. De Boever, M.B. Blaschko, Artery–vein segmentation in fundus images using a fully convolutional network, Comput. Med. Imaging Graph. 76 (2019) 101636.

[38] W. Ma, S. Yu, K. Ma, J. Wang, X. Ding, Y. Zheng, Multi-task neural networks with spatial activation for retinal vessel segmentation and artery/vein classification, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part I 22, Springer, 2019, pp. 769–778.

[39] F. Girard, C. Kavalec, F. Cheriet, Joint segmentation and classification of retinal arteries/veins from fundus images, Artif. Intell. Med. 94 (2019) 96–109.

[40] A. Galdran, M. Meyer, P. Costa, A. Campilho, et al., Uncertainty-aware artery/vein classification on retinal images, in: 2019 IEEE 16th International Symposium on Biomedical Imaging, ISBI 2019, IEEE, 2019, pp. 556–560.
[41] Z. Wang, X. Jiang, J. Liu, K.-T. Cheng, X. Yang, Multi-task siamese network for retinal artery/vein separation via deep convolution along vessel, IEEE Trans. Med. Imaging 39 (9) (2020) 2904–2919.
[42] J. Morano, Á.S. Hervella, J. Novo, J. Rouco, Simultaneous segmentation and classification of the retinal arteries and veins from color fundus images, Artif. Intell. Med. 118 (2021) 102116.
[43] W. Chen, S. Yu, K. Ma, W. Ji, C. Bian, C. Chu, L. Shen, Y. Zheng, TW-GAN: Topology and width aware GAN for retinal artery/vein classification, Med. Image Anal. 77 (2022) 102340.
[44] A. Galdran, A. Anjos, J. Dolz, H. Chakor, H. Lombaert, I.B. Ayed, State-of-the-art retinal vessel segmentation with minimalistic models, Sci. Rep. 12 (1) (2022) 6174.
[45] R.A. Karlsson, S.H. Hardarson, Artery vein classification in fundus images using serially connected U-nets, Comput. Methods Programs Biomed. 216 (2022) 106650.
[46] J. Hu, L. Qiu, H. Wang, J. Zhang, Semi-supervised point consistency network for retinal artery/vein classification, Comput. Biol. Med. 168 (2024) 107633.
[47] M. Mokan, G. Gabrani, D. Relan, Pixel-wise classification of the whole retinal vasculature into arteries and veins using supervised learning, Biomed. Signal Process. Control. (ISSN: 1746-8094) 106 (2025) 107691, http://dx.doi.org/10.1016/j.bspc.2025.107691, URL https://www.sciencedirect.com/science/article/pii/S1746809425002022.
[48] C. Wan, J. Cheng, W. Yang, L. Chen, DBMAE-net: A dual branch multi-scale feature adaptive extraction network for retinal arteriovenous vessel segmentation, Biomed. Signal Process. Control. 104 (2025) 107619.
[49] X. Chen, L. Niu, S. Guo, Efficient retinal artery/vein classification with dense color-invariant feature learning, Neural Comput. Appl. 37 (3) (2025) 1255–1270.
[50] M.E. Hoque, K. Kipli, T.M.A. Zulcaffle, R. Sapawi, A. Joseph, W.A.W.Z. Abidin, S.K. Sahari, Feature extraction method of retinal vessel diameter, in: 2018 IEEE-EMBS Conference on Biomedical Engineering and Sciences, IECBES, IEEE, 2018, pp. 279–283.
[51] A.W. Vaes, M.A. Spruit, N. Goswami, J. Theunis, F.M. Franssen, P. De Boever, Analysis of retinal blood vessel diameters in patients with COPD undergoing a pulmonary rehabilitation program, Microvasc. Res. 139 (2022) 104238.
[52] M. Niemeijer, X. Xu, A.V. Dumitrescu, P. Gupta, B. Van Ginneken, J.C. Folk, M.D. Abramoff, Automated measurement of the arteriolar-to-venular width ratio in digital color fundus photographs, IEEE Trans. Med. Imaging 30 (11) (2011) 1941–1950.
[53] D.K. Kumar, B. Aliahmad, H. Hao, Retinal vessel diameter measurement using unsupervised linear discriminant analysis, Int. Sch. Res. Not. 2012 (2012).
[54] J. Lowell, A. Hunter, D. Steel, A. Basu, R. Ryder, R.L. Kennedy, Measurement of retinal vessel widths from fundus images based on 2-d modeling, IEEE Trans. Med. Imaging 23 (10) (2004) 1196–1204.
[55] X.W. Gao, A. Bharath, A. Stanton, A. Hughes, N. Chapman, S. Thom, Quantification and characterisation of arteries in retinal images, Comput. Methods Programs Biomed. 63 (2) (2000) 133–146.
[56] A. Budai, R. Bock, A. Maier, J. Hornegger, G. Michelson, et al., Robust vessel segmentation in fundus images, Int. J. Biomed. Imaging 2013 (2013).
[57] Q. Hu, M.D. Abràmoff, M.K. Garvin, Automated separation of binary overlapping trees in low-contrast color retinal images, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013: 16th International Conference, Nagoya, Japan, September 22-26, 2013, Proceedings, Part II 16, Springer, 2013, pp. 436–443.
[58] J. Staal, M.D. Abràmoff, M. Niemeijer, M.A. Viergever, B. Van Ginneken, Ridge-based vessel segmentation in color images of the retina, IEEE Trans. Med. Imaging 23 (4) (2004) 501–509.
[59] K. Jin, X. Huang, J. Zhou, Y. Li, Y. Yan, Y. Sun, Q. Zhang, Y. Wang, J. Ye, Fives: A fundus image dataset for artificial intelligence based vessel segmentation, Sci. Data 9 (1) (2022) 475.
[60] E. Bullitt, G. Gerig, S.M. Pizer, W. Lin, S.R. Aylward, Measuring tortuosity of the intracerebral vasculature from MRA images, IEEE Trans. Med. Imaging 22 (9) (2003) 1163–1171.
[61] D.-Y. Yu, K.Y. Paula, S.J. Cringle, M.H. Kang, E.-N. Su, Functional and morphological characteristics of the retinal and choroidal vasculature, Prog. Retin. Eye Res. 40 (2014) 53–93.
[62] C.Y.-l. Cheung, C. Sabanayagam, A.K.-p. Law, N. Kumari, D.S.-w. Ting, G. Tan, P. Mitchell, C.Y. Cheng, T.Y. Wong, Retinal vascular geometry and 6 year incidence and progression of diabetic retinopathy, Diabetologia 60 (2017) 1770–1781.