Arteriosclerosis, Thrombosis, and Vascular Biology

BASIC SCIENCES

3DVascNet: An Automated Software for Segmentation and Quantification of Mouse Vascular Networks in 3D

Hemaxi Narotamo, Margarida Silveira, Cláudio A. Franco

BACKGROUND: Analysis of vascular networks is an essential step to unravel the mechanisms regulating the physiological and pathological organization of blood vessels. So far, most of the analyses are performed using 2-dimensional projections of 3-dimensional (3D) networks, a strategy that has several obvious shortcomings. For instance, it does not capture the true geometry of the vasculature and generates artifacts on vessel connectivity. These limitations are accepted in the field because manual analysis of 3D vascular networks is a laborious and complex process that is often prohibitive for large volumes.

METHODS: To overcome these issues, we developed 3DVascNet, a deep learning–based software for automated segmentation and quantification of 3D retinal vascular networks. 3DVascNet performs segmentation based on a deep learning model, and it quantifies vascular morphometric parameters such as vessel density, branch length, vessel radius, and branching point density. We tested the performance of 3DVascNet using a large data set of 3D microscopy images of mouse retinal blood vessels.

RESULTS: We demonstrated that 3DVascNet efficiently segments vascular networks in 3D and that vascular morphometric parameters capture phenotypes detected by using manual segmentation and quantification in 2 dimension. In addition, we showed that, despite being trained on retinal images, 3DVascNet has high generalization capability and successfully segments images originating from other data sets and organs.

CONCLUSIONS: Overall, we present 3DVascNet, a freely available software that includes a user-friendly graphical interface for researchers with no programming experience, which will greatly facilitate the ability to study vascular networks in 3D in health and disease. Moreover, the source code of 3DVascNet is publicly available, thus it can be easily extended for the analysis of other 3D vascular networks by other users.

GRAPHIC ABSTRACT: A graphic abstract is available for this article.

Key Words: deep learning ■ phenotype ■ radius ■ retinal vessels ■ software

Blood vessel organization and function are essential for embryogenesis and homeostasis, and their dysregulation is associated with several diseases, including cancer, cardiovascular diseases, neurodegenerative disorders, and aging. Blood vessels form an intricate and complex network of vessels that irrigate all tissues. The 3-dimensional (3D) hierarchical architecture of vessels is relevant for their correct function and is stereotypical to each organ.[1–3] Yet, quantitative information on vascular parameters is mostly extracted from maximum-intensity projections (MIPs) of stacks of images, that is, quantifications are performed in 2-dimensional (2D) images.[4–9]

Researchers have adopted this conscious limitation because manual analysis of 3D vascular networks is a laborious and complex process that is often prohibitive for large volumes. However, 2D analysis brings obvious

---

Correspondence to: Cláudio A. Franco, PhD, Católica Biomedical Research Centre, Rua da Quinta Grande 6, 2780-156 Oeiras, Portugal. Email cfranco@ucp.pt

Supplemental Material is available at https://www.ahajournals.org/doi/suppl/10.1161/ATVBAHA.124.320672.

For Sources of Funding and Disclosures, see pages 1599–1600.

Arterioscler Thromb Vasc Biol is available at www.ahajournals.org/journal/atvb

Nonstandard Abbreviations and Acronyms

2D	2-dimensional
3D	3-dimensional
Ang II	angiotensin II
CD31	cluster of differentiation 31
clDice	centerline Dice
DC	Dice coefficient
DL	deep learning
GUI	graphical user interface
ICAM2	intercellular adhesion molecule 2
MIP	maximum-intensity projection
NMI	normalized mutual information
ROI	region of interest
VEGF	vascular endothelial growth factor

Highlights

- 3DVascNet is a fully automated open-source tool to segment and quantify 3D vascular networks.
- Segmentation is performed using a deep learning model that does not require 3D manual annotation for training.
- Quantifications based on 2D projections of 3DVascNet masks and 2D manually annotated masks do not present statistically significant differences.
- 3DVascNet's user-friendly graphical interface can be easily used by researchers with no programming knowledge.
- 3DVascNet will be a useful tool for vascular biologists, aiding the study of vascular networks in health and disease.

shortcomings. For instance, analyses based on 2D MIPs from 3D images neglect the 3D anatomic structure of the vasculature, thereby missing important 3D biological features and not capturing the full complexity of phenotypes involved in physiology and pathology. Recently, researchers have started to use deep learning (DL) methods with automated and semiautomated image processing pipelines to segment vascular networks in 3D. For instance, Todorov et al[10] presented VesSAP, an automated tool for 3D segmentation and quantification of the mouse brain vasculature. In VesSAP, segmentation is performed using a 3D convolutional neural network. Nevertheless, this method and most of the DL models that have been developed for vessel segmentation require manually annotated data sets for training.[11–13] However, manual segmentation is a time-consuming process, and, particularly for 3D blood vessels, it is often prohibitive, which hinders significantly the development of such methods.

To overcome these limitations, we have developed 3DVascNet, a DL-based software for automated segmentation and quantification of 3D vascular networks, which does not require pairs of images and manually annotated ground-truth segmentation masks. To develop 3DVascNet, we focused on retinal blood vessels. Dysfunction of this stereotypical vascular network is associated with ocular diseases, such as diabetic retinopathy, age-related macular degeneration, and choroidal neovascularization,[11,14] and it is one of the reference models to investigate angiogenesis.[14] Hence, over the years, significant attention has been paid to retinal blood vessel formation and remodeling.[11,14]

Several alterations in the morphology of the blood vessels provide valuable information for the diagnosis of retinal diseases. Various geometric features are typically used to describe the structure and morphology of the blood vessels: diameter, length, area, perimeter, tortuosity, number of branches, and branching points and angles.[13,15] On the one hand, for clinicians, the analysis and interpretation of these features are important for the diagnosis of several pathological conditions.[11–13] On the other hand, for researchers, the vessels' features are relevant to understand the mechanisms underlying physiological and pathological retinal vasculature formation,[14] which is fundamental to develop new therapeutic approaches.[16]

The proposed 3DVascNet software enables the analysis of blood vessel networks in 3D, which will greatly facilitate the ability to quantify and study vascular networks in health and disease. 3DVascNet is based on a CycleGAN model,[17] a generative DL model designed to translate images from domain A into domain B. In this work, domain A consists of 3D microscopy images of retinal vessels, and domain B corresponds to 3D segmentation masks of vessels. To the best of our knowledge, this is the first work proposing an automated approach for segmentation and quantification of 3D retinal blood vessels, which does not require pairs of images and corresponding ground-truth masks for training. The source code of 3DVascNet is freely available; thus it can be easily extended to analyze other 3D vasculatures. Moreover, we provide a graphical user interface (GUI) suitable for researchers without a solid knowledge of programming.

The main contributions of this work are the following:

- We propose a DL-based approach for 3D retinal vessel segmentation in microscopy images that does not require paired 3D images and 3D masks for training;
- 3DVascNet, available as source code and GUI, can be used for automated and accurate quantification of the 3D retinal vasculature;
- We demonstrate the robustness of our approach by testing it on 4 different data sets of mouse retinas;
- We release 4 data sets containing 3D microscopy images of retinal blood vessels and corresponding 2D and 3D masks.

METHODS

Availability of Data

All data and materials have been made publicly available at the Github and HuggingFace repositories and can be accessed at https://github.com/HemaxiN/3DVascNet and https://huggingface.co/datasets/Hemaxi/3DVesselSegmentation. Please see the Major Resources Table in the Supplemental Material.

Data Sets

In brief, the data set used in this work comprises 3D images of mouse retinas acquired and analyzed in a previous work.[6] This data set comprises 21 3D microscopy images of mouse retinal vessels, the corresponding 2D MIP images, and the corresponding 2D and 3D masks. Both the 2D and 3D masks were created by Barbacena et al[6] using the PolNet method.[18] In this method, first the 2D masks are obtained from the 2D MIP images, then the 3D masks are computed from the 2D masks. Specifically, the 2D masks are obtained by applying a traditional 2D segmentation workflow to the MIP images, which includes thresholding, outliers and unconnected object removal, morphological operations (binary erosion and dilation), and manual correction.

The 3D masks are obtained from the 2D masks using a set of assumptions.[18] Briefly, the skeletons and radii of each vessel segment are computed from the 2D masks. Retinal vessel skeleton is the representation of the vascular vessels as a 1-pixel-wide line containing pixels that are equidistant from the vessel boundaries, which is also denoted as the vessel centerline. Afterward, the 3D masks are generated considering that the vessels' 3D structure can be approximated as tubular segments, where each segment is cylindrical with a circular cross-section. The radius of each vessel segment in 3D is presumed to be equal to the radius computed from the corresponding 2D projection. Since these PolNet's 3D masks are not obtained directly from the 3D images, there is no overlap between these two across the *Z* axis. The only paired data are the 2D MIP images and the 2D masks. The 3D masks are a coarse representation of the vessels in the 3D images.

The PolNet method was only applied to retinas until postnatal day 7, because until approximately this developmental stage, the retina has a single and coplanar network of blood vessels, the superficial layer, in the nerve fiber layer. After postnatal day 7, the vessels in the single layer begin to sprout vertically, resulting in the formation of a second layer, approximately parallel to the first layer, and then a third layer also approximately parallel to the first and second layers. The PolNet method is not applicable to retinas beyond postnatal day 7 because for these retinas, the corresponding 2D MIP images would contain overlap between vessels of different plexuses/layers, which would lead to incorrect estimation of branches and branching points and consequently would lead to the generation of a 3D mask with an incorrect topology. This limitation was already identified in the PolNet paper.[18]

Table 1 describes data set 1, where the prefix in each image name (first column) denotes a specific treatment provided to the retinas as described in the Table's caption. Details about the staining procedures can be found in the Method Details section, in the In Vivo Mouse Treatments subsection in the study by Barbacena et al.[6] We used a single image from each condition to train the 3DVascNet model; the remaining images were used as test data to evaluate the model's performance.

Furthermore, to test the generalization capability of our segmentation model, we evaluated its performance using new data that differs from the training data set in terms of image acquisition conditions (different users and different microscopes) and retinal treatments. Specifically, we considered 3 other data sets (data sets 2–4) containing 28 3D images and 2D masks (Table S1). These data sets contain images of retinas that have undergone treatments other than those listed in Table 1 and were acquired by different users from those who obtained data set 1. Additionally, 7 of these 28 images were acquired in another laboratory and consequently using a different microscope from the one used to obtain the training images. All data sets are publicly available at https://huggingface.co/datasets/Hemaxi/3DVesselSegmentation.

3DVascNet's Pipeline

The training and testing pipelines of 3DVascNet are depicted schematically in Figure 1A. First, a 3D CycleGAN model is trained using image volumes and volumes extracted from PolNet masks (Figure 1A, top). Thereafter, the testing pipeline is applied to test images; it consists of 6 main modules

Table 1. Description of Data Set 1 Used in This Work for Training and Testing

Image	Size (X×Y×Z)	Train/test split
PBS_ret1	5409×2979×31	Train
PBS_ret2	3375×4599×56	Test
PBS_ret3	6237×3384×38	Test
PBS_ret4	5013×2979×43	Test
PBS_ret5	4838×5561×39	Test
PBS_ret6	4864×4864×88	Test
PBS_ret7	6144×8192×20	Test
sFLT1_ret1	4617×5418×35	Train
sFLT1_ret2	4365×3978×73	Test
sFLT1_ret3	4365×3600×66	Test
VEGF_ret1	2979×4599×42	Train
VEGF_ret2	5319×4878×61	Test
VEGF_ret3	5427×4194×46	Test
VEGF_ret4	5013×4194×38	Test
VEGF_ret5	4365×3978×40	Test
VEGF_ret6	4743×4365×61	Test
VEGF_ret7	4365×4743×42	Test
AngII_ret1	6273×3780×33	Train
AngII_ret2	4851×5571×42	Test
AngII_ret3	4864×3584×48	Test
AngII_ret4	4864×5120×40	Test

The prefix in each image name (first column) represents the treatment that was applied to the retina. Retinas were collected from animals treated with (1) intraperitoneal injection of PBS, which works as a control condition; (2) intraocular injection of sFLT1, a protein that blocks VEGFA, which acts as an antiangiogenic molecule; (3) intraocular injection of VEGFA, a proangiogenic molecule; and (4) intraperitoneal injection of Ang II, a vasoconstrictor peptide, which increases blood pressure. These images were previously collected and analyzed in the study by Barbacena et al.[6] Ang II indicates angiotensin II; sFLT1, soluble fms-like tyrosine kinase 1; and VEGFA, vascular endothelial growth factor A.

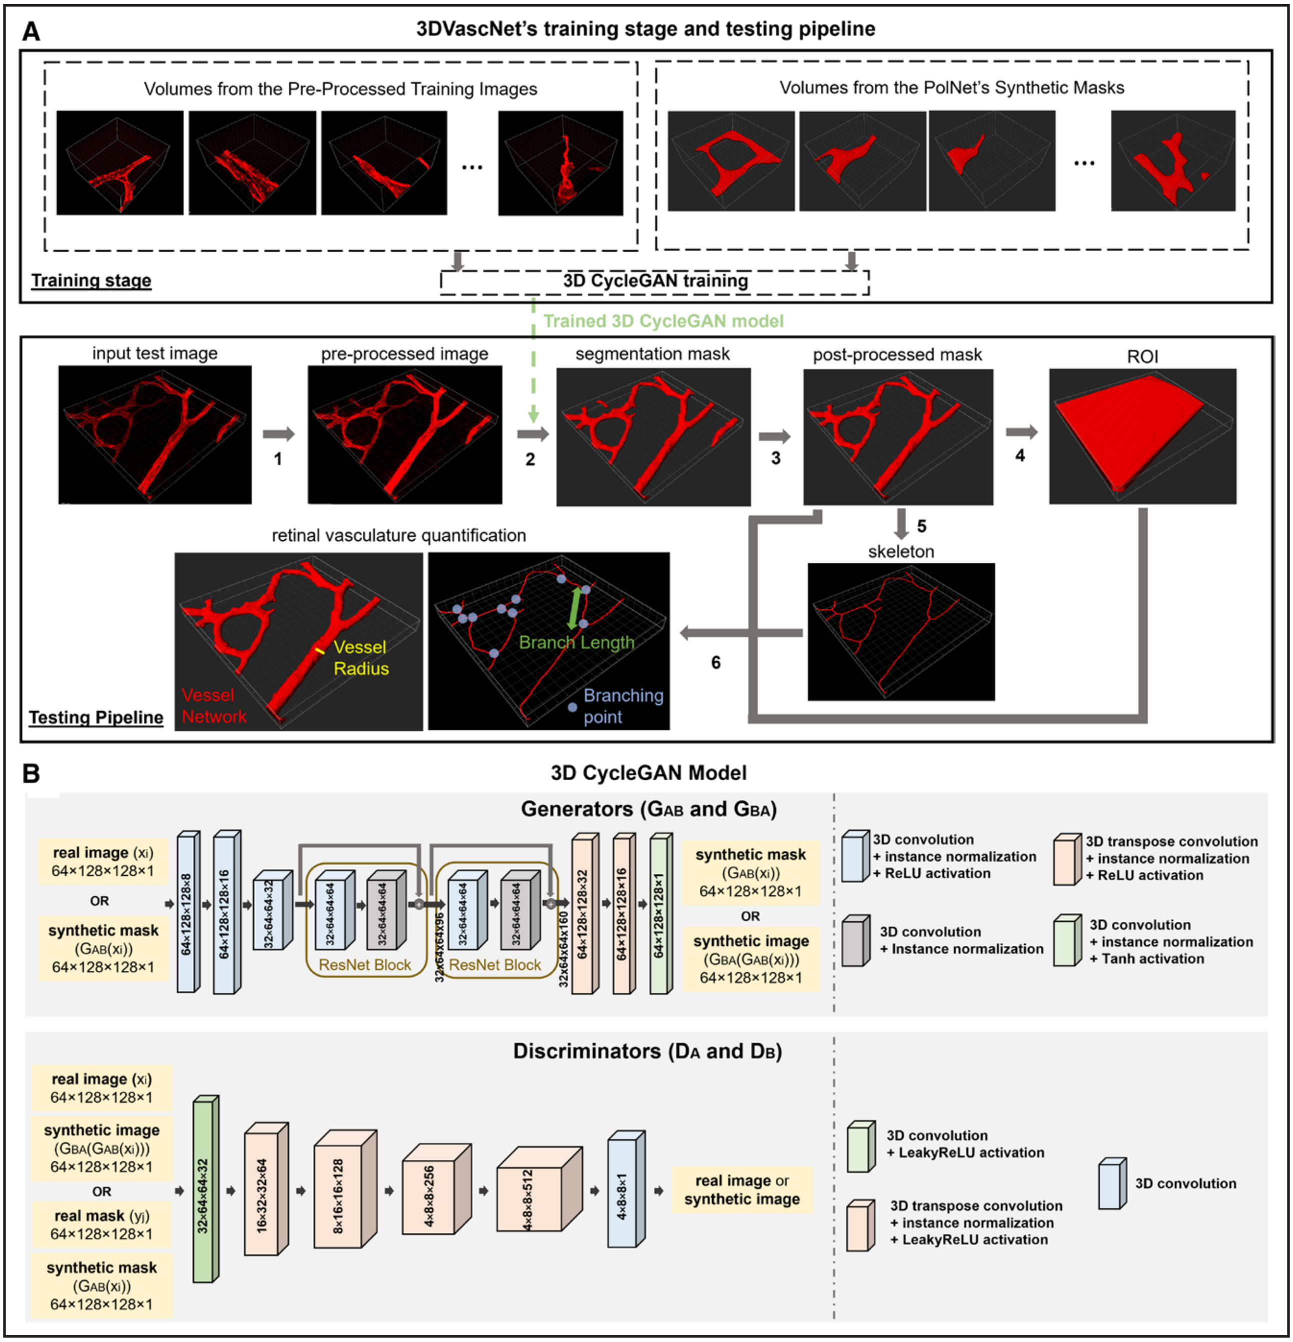

Figure 1. 3DVascNet's pipeline and 3-dimensional (3D) CycleGAN model architecture.
A, Overview of the 3DVascNet training and testing stages. **Top**, Training a 3D CycleGAN model using image volumes and synthetic mask volumes. **Bottom**, 3DVascNet testing pipeline that consists of 6 main blocks: (1) image preprocessing, (2) image segmentation based on the trained 3D CycleGAN model, (3) mask postprocessing, (4) automatic computation of the region of interest (ROI), (5) mask skeletonization, (6) feature extraction (vessel density, vessel radius, branch length, branching points density). **B**, Architecture of the 3D CycleGAN model for retinal vessel segmentation in microscopy images. It is composed of 2 generators (**top**) and 2 discriminators (**bottom**). ResNet denotes the residual network, which includes skip connections to improve training of deep learning models, and ReLU (rectified linear unit) is an activation function. The numbers in the boxes denote the size of the feature maps resulting from the operation performed by the box.

(Figure 1A, bottom): (1) image preprocessing to enhance the vessels, (2) image segmentation based on the trained 3D CycleGAN model, (3) mask postprocessing, (4) automated region of interest (ROI) computation, (5) mask skeletonization, and (6) feature extraction based on the masks and skeletons. More information on individual blocks of these pipelines is provided in the following subsections.

Pre-Processing

Several preprocessing methods have been proposed for intensity inhomogeneity correction, image normalization, and vessel enhancement to improve retinal vessel segmentation.[11,13,19] Here, we performed normalization for postnatal day 6 retinas based on the 1st and 99th percentiles of the image intensity values, similar to the studies by Montoya-Zegarra et al[8] and

Todorov et al.[10] In this way, the voxels in the 1% to 99% range of intensity values are linearly mapped into the (0,1) interval. Moreover, the bottom and top 1% of voxels are clipped at 0 and 1, respectively. Figure S1A shows an example of a patch extracted from the original microscopy image and the corresponding patch extracted from the preprocessed image.

We noticed that while percentile normalization with values of 1 and 99 is effective for single-layer vasculature (postnatal day 6 retinas), it was not optimal for adult retinas, where we found substantial intensity variations between different vessel layers. We have found that using percentiles of 0.5 and 99.5 yields better results. Thus, we recommend adapting this normalization step, especially when using complex vasculatures.

3D Blood Vessel Segmentation

In this section, we present the proposed DL-based approach for 3D blood vessel segmentation and its training and testing details. To overcome the problem of manually annotating 3D masks for training, in this work, we propose a 3D CycleGAN model that only requires examples of images from the source domain and target domain for training. It does not require pairs of 3D images and masks. The source domain consists of 3D microscopy images, and the target domain comprises 3D masks that were created based on 2D masks of retinal vessels using the approach proposed in the study by Bernabeu et al,[18] as explained in the Datasets subsection. Thus, manual annotation of 3D blood vessels is not needed in our approach.

3D CycleGAN Model

The proposed 3D CycleGAN model is an extension to 3D of the 2D CycleGAN model developed for unpaired image-to-image translation.[17] To implement the 3D CycleGAN model based on the 2D CycleGAN model,[17] we replaced the 2D layers (convolutional and transposed convolutional) by 3D layers. The 3D CycleGAN model is composed of 2 generators (G_{AB} and G_{BA}) and 2 discriminators (D_A and D_B), where G_{AB} maps volumes from domain A (source domain) into domain B (target domain), G_{BA} translates volumes from domain B into domain A, and D_A and D_B classify, respectively, volumes from domains A and B as real or synthetic. The architecture of the generators and discriminators is depicted in Figure 1B.

Each generator contains 3D convolutional, instance normalization, and ReLU activation layers followed by a downsampling layer that reduces the size of the feature maps. Thereafter, it contains residual blocks that are composed of 3D convolutions, instance normalization, and ReLU activation. Finally, an upsampling layer increases the size of the feature maps, and it is followed by 3D transposed convolution, instance normalization, and ReLU activation layers, except the activation function of the last layer, which is a Tanh activation function as proposed in the study by Zhu et al.[17] The discriminator is a convolutional neural network designed for image classification. It takes as input a real or synthetic volume and outputs the probability of this volume being real or synthetic. The discriminator is composed of 3D convolutional, instance normalization, and LeakyReLU activation layers (Figure 1B).

Considering an unpaired data set containing volumes of microscopy images ($\{x_i\}_{i=0}^{N-1}$, $x_i \in$ domain A) and volumes of masks ($\{y_j\}_{j=0}^{M-1}$, $y_j \in$ domain B), the mapping $G_{AB}: \mathrm{A} \to \mathrm{B}$ receives an image $x_i \in$ domain A and learns to generate a mask $G_{AB}(x_i)$ that is indistinguishable from mask $y_j \in$ domain B. It learns this mapping by adversarial training as proposed in the study by Isola et al.[20] That is, on the one hand, G_{AB} learns to generate synthetic masks, on the other hand, D_B tries to distinguish real masks (y_j) from synthetic ones ($G_{AB}(x_i)$), and the goal of G_{AB} is to fool the discriminator D_B by generating synthetic masks similar to the real ones. Nevertheless, this does not ensure that x_i and $G_{AB}(x_i)$ will be paired. Therefore, structure is added to the 3D CycleGAN to encourage that if we transform an image (x_i) into a mask ($G_{AB}(x_i)$) and then transform it back into an image ($G_{BA}(G_{AB}(x_i))$), we get an image similar to x_i. This is expressed as $G_{BA}(G_{AB}(x_i)) \approx x_i$, where $G_{BA}: B \to A$ is the mapping from the masks set into the images set. To add this structure, the 3D CycleGAN has 2 training paths: forward ($x_i \to G_{AB}(x_i) \to G_{BA}(G_{AB}(x_i)) \approx x_i$) and backward ($y_j \to G_{BA}(y_j) \to G_{AB}\left(G_{BA}(y_j)\right) \approx y_j$), as represented in Figure S1B (top).

Thus, mappings $G_{AB}: A \to B$ and $G_{BA}: B \to A$ along with their respective generators, D_B and D_A, are trained through adversarial training. More specifically, the training process is cyclical where, at each iteration, $G_{BA}: B \to A$ is updated, followed by training D_A using the generated samples, then $G_{AB}: A \to B$ is updated, followed by updating D_B using the generated samples. Moreover, the addition of the cycle consistency loss encourages $G_{AB}\left(G_{BA}(y_j)\right) \approx y_j$ and $G_{BA}(G_{AB}(x_i)) \approx x_i$. Finally, an identity term is added to preserve the color. Hence, the objective function of the 3D CycleGAN model contains 4 terms:

$$\begin{aligned} L(G_{AB}, G_{BA}, D_A, D_B) &= L_{GAN}(G_{AB}, D_B, x, y) \\ &+ L_{GAN}(G_{BA}, D_A, y, x) + \lambda_1 L_{Cyc}(G_{AB}, G_{BA}) \\ &+ \lambda_2 L_{identity}(G_{AB}, G_{BA}), \end{aligned} \tag{1}$$

where λ_1 and λ_2 are constants. The first 2 terms are the adversarial losses:

$$\begin{aligned} L_{GAN}(G_W,\ D_W,\ a, b) &= \mathbb{E}_b\left[(D_W(b) - 1)^2\right] \\ &+ \mathbb{E}_a\left[(D_W(G_w(a)))^2\right] \end{aligned} \tag{2}$$

where $\mathbb{E}_a$ and $\mathbb{E}_b$ represent the expectation of the log-likelihood of a sample being sampled from the probability distribution of a and b, respectively. The goal of G_W is to maximize this objective, and the goal of D_W is to minimize it: $\max_{G_W} \min_{D_W} L_{GAN}(G_W, D_W, a, b)$. In this work, the loss shown in Equation 2 is applied to the pairs $\{G_{AB}, D_B\}$ and $\{G_{BA}, D_A\}$.

The third term is the cycle consistency loss:

$$\begin{aligned} & L_{Cyc}(G_{AB},\ G_{BA}) = \\ & \mathbb{E}_x\left[\|G_{BA}(G_{AB}(x)) - x\|_1\right] + \\ & \mathbb{E}_y\left[\|G_{AB}(G_{BA}(y)) - y\|_1\right]. \end{aligned} \tag{3}$$

The last term is the identity loss:

$$L_{identity}(G_{AB},\ G_{BA}) = \mathbb{E}_x\left[\|G_{BA}(x) - x\|_1\right] + \mathbb{E}_y\left[\|G_{AB}(y) - y\|_1\right]. \tag{4}$$

Our implementation of the 3D CycleGAN model is made publicly available at https://github.com/HemaxiN/3DVascNet, and it is based on a 2D implementation.[21]

Training

To train the 3D CycleGAN model, we extracted nonoverlapping volumes of size 128 × 128 × 64 from the 3D images and 3D masks. Padding was performed before extracting the volumes, and the new dimensions ($\tilde{s}izex \times \tilde{s}izey \times \tilde{s}izez$) of the images/masks were computed as follows:

1. $\tilde{s}izex = \mathrm{int}\,(\mathrm{sizex}/\mathrm{px} + 1) \times \mathrm{px}$
2. $\tilde{s}izey = \mathrm{int}\,(\mathrm{sizey}/\mathrm{py} + 1) \times \mathrm{py}$
3. $\tilde{s}izez \quad = \quad 64$

where *sizex* and *sizey* denote the original sizes of the image/mask along the x and y dimensions, $int(\cdot)$ converts the specified value into the largest integer not greater than the value, and px and py denote the patch size along the x and y axes, respectively. In this way, the image size is adapted to be a multiple of the patch size along x and y. For the 3D masks, we performed zero padding. The original and preprocessed 3D images were padded, respectively, with small random values in the intervals (1,25) and (1,50), since these are the typical voxel intensity values in the background of these images.

After performing padding and extracting volumes from the images and masks, we obtained 4987 image volumes and 4987 mask volumes for training, which were normalized to present values ranging from −1 to 1. Based on our experiments, we concluded that training the model for a large number of epochs deteriorated the results. For instance, after 11 epochs, the output of the model was a tensor of zeros (Figure S2A). The best models were obtained between epochs 1 and 5. We trained the 3D CycleGAN model from scratch for 15 epochs, using Adam optimizer (learning rate of 1×10^{-4} and $\beta_1 = 0.5$)[22] and setting $\lambda_1 = 10$ and $\lambda_2 = 5$ as suggested in the study by Zhu et al[17] and defined in Equation 1. The batch size was set to 1 due to memory limitations. Training the 3D CycleGAN model with this configuration took ≈36 hours. The best epoch was selected based on visual inspection of the segmentation masks at each epoch and considering an unsupervised evaluation metric that is described in Evaluation Metrics (Figure S2B).

Testing

For testing, we extracted volumes of size $128 \times 128 \times 64$ with an overlap of 64 voxels along the x and y directions and 32 voxels along the z direction. These volumes were extracted from the padded images. For the test images, padding was performed as explained above, except for images with >64 slices. For these images, the new dimension along z was computed as follows: $\tilde{s}izez = int\,(sizez/pz + 1) \times pz$, where *sizez* is the original image size across z and pz is the patch size along the z axis. The volumes were then fed to the trained generator $G_{AB} : A \rightarrow B$ (Figure S1B, bottom), which outputs volumes (out) with values ranging from −1 to 1. First, the following operation was performed to ensure that the output volume values are within the range $[0, 1] : \widetilde{out} = \frac{(out+1)}{2}$. Afterward, a threshold of 0.5 was applied to out: voxels with values <0.5 were set to 0, indicating background, while all others were set to 1, representing the vessels. To incorporate patches with overlapping regions, we used a voting mechanism based on 2 variables: N_{back} and N_{fore}. Here, $N_{back}\,[u, v, x]$ and $N_{fore}\,[u, v, x]$ count the occurrences of voxel (u, v, x) being classified as background and foreground, respectively. Thereafter, the voxel values in the output segmentation mask were computed as follows:

$$mask = N_{fore} \geq N_{back}. \tag{5}$$

Finally, we cropped the segmentation masks along x and y to match the size of the corresponding images. The test time for about 37 000 patches of size 128×128×64 is ≈1 hour.

Post-Processing

After performing vessel segmentation, we applied a simple postprocessing step, which aims at removing small vessel segments and bright background noise due to impurities or staining irregularities. We selected all objects with a volume >50 000 µm^3 and all objects with a volume between 1000 and 50 000 µm^3, which have convexity <0.2. All objects with a volume inferior to 1000 µm^3 were excluded. Convexity is computed as the ratio between the object's volume and the volume of the object's convex hull, which is the smallest convex region that fits the object. This parameter is useful for distinguishing small, isolated blobs from isolated blood vessels, which usually have convexity values above 0.2. In our training and testing data set, we used a more stringent threshold to only select the largest object, to match the threshold used in PolNet data sets (data sets 1–4). For all the other images and data sets, we used the default threshold mentioned above. Users may consider changing this postprocessing step and adjust it to their specific data set requirements.

Definition of an ROI

The segmentation masks typically contain the vascular structure and background regions. Thus, we need to define a ROI to extract features only from the vascular network. This allows for more targeted analysis and avoids unnecessary computations on irrelevant regions.

In this work, we automatically defined the 3D ROIs for each image. First, we obtained the 3D convex hull from each 3D segmentation mask. Thereafter, we computed the concave hull of each 2D segmentation mask[23] and applied some preprocessing and postprocessing steps (Figure S3A). The 2D masks were obtained from the 3D segmentation masks by performing an AND operation along the z direction. To compute the concave hull, first, the segmentation mask was resized to 100,100 to decrease the computational cost. Then, the coordinates of the points belonging to the vessel edges were used to compute the concave hull as proposed in the study by Moreira et al.[23] Thereafter, the resulting binary mask was resized to the original size, and 2 postprocessing steps were applied. First, regions outside the concave hull that were <5% of the largest region outside the concave hull were added to the ROI. Thereafter, median filter was applied to reduce staircase-like borders shown in Figure S3A. Lastly, to create the final 3D ROI, we performed an intersection of the 2D ROI with each slice along the z direction obtained from the 3D convex hull (Video S1).

Vascular biologists typically perform selective feature extraction to focus on specific areas of interest. First, they define several ROIs on the images; thereafter, they extract features from these ROIs to study vascular networks. While our automated approach computes a single ROI for the input image, we have an optional step in our GUI that allows users to manually select multiple ROIs.

Skeletonization

After obtaining the vessels' segmentation masks and ROIs, the next step consists in extracting the vessels' skeletons. Before computing the skeleton, we applied a 3D morphological closing operation with a binary ellipsoid of size 9×9×3 to fill small holes in the 3D segmentation masks.

A 3D segmentation mask and the corresponding retinal vessel skeleton are represented in Figure S3B. A branching point is a point connected to ≥3 vessel segments, and an end point is a point connected to one vessel segment (Figure S3B). A branch is a segment between 2 branching points or between an end point and a branching point.

In this study, we only have 2D ground-truth masks; thus we extracted 2D skeletons from the predicted and ground-truth masks to compare retinal vasculature quantification. The predicted 3D segmentation masks were projected into 2D using an AND operation along the *z* axis. To avoid false-positive centerlines, as proposed in the study by Todorov et al,[10] we applied erosion followed by dilation to the 2D masks using an ellipsoid structuring element of size 5×5. Then, the skeleton of each segmentation mask was computed inside the ROI using the method proposed by Lee et al.[24]

However, as explained in the introduction, 2D skeletons do not accurately represent the real structure of 3D vessels. Herein, they were computed only for comparison purposes. Our 3DVascNet software computes the 3D skeletons from the 3D segmentation masks using the method presented in the study by Lee et al,[24] to perform retinal vasculature quantification in 3D.

Retinal Vasculature Features

In this work, based on both the segmentation masks and their skeletons, we extracted 4 features to describe retinal blood vessels: vessel density (vascular volume/ROI volume), mean vessel radius, mean branch length, and branching points density.

For each image, the 4 abovementioned features were computed within the ROI:

- Vessel density (vascular area/ROI area): we compute the quantity $\frac{V_{v\in ROI}}{V_{ROI}} \times 100$, where V_{ROI} and $V_{v\in ROI}$ denote the volume (number of voxels) of the ROI and of the vessels inside the ROI, respectively.
- Mean vessel radius: briefly, for each branch in the skeleton, its direction was computed, and the vessel cross-section orthogonal to that direction was determined for each point (p) in the branch. The vessel contour was extracted from the cross-section, which is composed of N points ($c_0, c_1, \ldots, c_{N-1}$). The mean vessel radius for each point in the branch was computed as the mean of the Euclidean distances between each contour point and the branch point: $\frac{1}{N}\sum_{c=0}^{N-1}\|c_n - p\|_2$. The mean vessel radius corresponds to the mean of the radii computed for all points in all branches.
- Branching points density: we compute $\frac{\#BP}{V_{ROI}}$, where $\#\ BP$ denotes the number of branching points.
- Mean branch length: it represents the mean of the individual branches' lengths. That is, it is the mean of the sum of the Euclidean distances between consecutive voxels along each branch in the skeleton.

The length of the branches and the number of bifurcation points were computed using the Skan tool.[25] The voxel spacing was taken into consideration to compute the branch lengths and vessel radii in µm. These features were also computed in 2D for comparison purposes.

Compared Approaches

The approaches that we compared with are the following:

- Otsu: it is an easy and fast method to segment images. It selects a threshold based on the histogram of image intensities. The threshold is selected to maximize interclass variance.[13,26,27] For a bimodal histogram, the selected threshold separates the 2 histogram peaks.
- Chan-Vese[28]: it is an iterative algorithm that belongs to a family of methods for image segmentation named geometric active contours or snakes. Chan-Vese works with the assumption that the foreground and background pixel intensity values follow 2 gaussian distributions. Thus, it presents good results when the background and foreground regions have different means. Segmentation is obtained by minimizing an energy function, which is the weighted sum of 4 terms (Equation 6): (1) the sum of the squared differences between the intensities of the voxels in the segmented region and the mean intensity of this region, (2) the sum of the squared differences between the intensities of the voxels in the background region and the mean intensity of this region, (3) a term dependent on the volume of the segmented region, and (4) another term dependent on the length of the segmented region's boundary:

$$F(c_1, c_2, C) = \lambda_1 \int_{inside(C)} |u_0(q,r,s) - c_1|^2 dq\,dr\,ds + \lambda_2 \int_{outside(C)} |u_0(q,r,s) - c_2|^2 dq\,dr\,ds + \nu \cdot Volume(inside(C)) + \mu \cdot Length(C) \quad (6)$$

where u_0 denotes the image to be segmented; C, the curve that encloses the foreground region, c_1 and c_2 are the averages of the voxel intensity values inside and outside C, respectively; and $\nu \geq 0$, $\mu \geq 0$, $\lambda_1 > 0$, and $\lambda_2 > 0$ are the weights applied to each term.

To test this algorithm in our data set, we used a publicly available implementation: https://github.com/pmneila/morphsnakes.

Evaluation Metrics

Several supervised evaluation metrics have been proposed for the quantitative assessment of retinal vessel segmentation approaches: Dice coefficient (DC),[29] accuracy, sensitivity, and specificity, among others. In these metrics, the segmentation masks obtained with a given method are compared with the respective ground-truth masks, which are typically manually annotated by experts.[11–13,19,30] Herein, we used the DC metric[29]:

$$DC(VGt, VPred) = 2\frac{|VGt \cap VPred|}{|VGt| + |VPred|} \quad (7)$$

where $|VGt \cap VPred|$ denotes the number of pixels/voxels, classified as vessels, belonging to the intersection between the ground-truth masks (VGt) and predicted masks (VPred), and $|VGt|$ and $|VPred|$ denote the number of pixels/voxels classified as vessels in VGt and VPred, respectively. VGt and VPred are binary masks, where 1 denotes vessels and 0 denotes background.

Moreover, since we are interested in studying the topology of the blood vessels, we also considered a recently proposed metric, the centerline Dice (clDice)[31]:

$$clDice(VPred,\ VGt) = 2 \times \frac{T_{Prec}(SPred, VGt) \times T_{sens}(SGt, VPred)}{T_{Prec}(SPred, VGt) + T_{sens}(SGt, VPred)} \quad (8)$$

where $T_{Prec}(SPred, VGt) = \frac{|SPred\ \cap\ VGt|}{|SPred|}$, $T_{sens}(SGt, VPred) = \frac{|SGt\ \cap\ VPred|}{|SGt|}$, and $SPred$ and SGt are the skeletons of $VPred$ and VGt, respectively. SGt and $SPred$

are binary volumes, where 1 denotes vessels' skeleton and 0 background.

However, these supervised metrics require manual segmentation masks that are difficult to obtain. Besides, they may be subjective due to the interexpert variability associated to manual segmentation.[32] Herein, since we already had the 2D masks that were generated in a previous study,[6] we used those to compute the DC and clDice values to evaluate the performance of the compared segmentation approaches.

Unsupervised evaluation metrics do not require ground-truth segmentation masks. Instead, these metrics are computed based on the comparison between each segmentation mask generated by a given model and the corresponding original input image. Zhang et al[32] evaluated the performance of several unsupervised metrics, such as entropy, texture, and color error. They computed the accuracy for each unsupervised metric, which represents the number of times the metric gives a better score to the mask annotated by a human evaluator instead of the mask generated by a segmentation model.[32]

Unsupervised metrics can be computed for any segmented image and are objective.[32] Nevertheless, these metrics have disadvantages, for instance, given 2 masks A (optimal segmentation) and B (noisy segmentation), an unsupervised metric does not always give a better segmentation score to mask A. Zhang et al[32] obtained accuracy values between 1.0% and 82.1% for the studied unsupervised evaluation metrics.

In this work, to select the best epoch of the 3D CycleGAN model, we propose the normalized mutual information (NMI)[33] as an unsupervised segmentation evaluation metric because it is easier to compute than the metrics studied in the study by Zhang et al.[32] NMI is a widely used metric to measure the alignment between 2 images of different modalities, without requiring the signals to be identical in these 2 images. It measures the predictability of the signal in one image given the signal in another image. This metric is appropriate to measure the level of similarity between the original microscopy images, with values in the interval (0,255) and the corresponding binary segmentation masks with zeros and ones because it assesses how well the segmentation captures the vascular structures present in the image (Figure S2). NMI is computed as follows:

$$NMI\,(I, M) = \frac{H\,(I) + H\,(M)}{H\,(I, M)} - 1 \tag{9}$$

where

$$H\,(N) = -\sum_{n \in N} p\,(n) \times log\,(\,p\,(n)) \tag{10}$$

is the marginal entropy; it represents the average information for an image (N) containing the set of values n ($n \in N$), and $p\,(n)$ denotes the marginal probabilities of individual values,[33] and

$$H\,(M, N) = -\sum_{n \in N}\sum_{m \in M} p\,(m, n) \times log\,(\,p\,(m, n)) \tag{11}$$

is the joint entropy of images M and N, and the joint probability $p\,(m, n)$ determines how frequently pairs of values occur together.[33]

3DVascNet Performance on 3D Ground Truth

We manually generated 3D ground-truth for a small subset of our data set using a 2-stage strategy. In the first step, the image intensity was manually adjusted to enhance the vessels, and thresholding was applied to create a coarse 3D mask. Afterward, since we have 2D ground-truth data, we performed intersection of each 2D slice of the coarse 3D mask with the 2D ground-truth mask. Thereafter, the masks were manually corrected, which included tasks such as hole filling, boundary adjustment, and inclusion/exclusion of regions, by a junior researcher and a senior experienced biologist, to ensure accuracy.

The ground truth was created for 9 small patches of size 600×600×64 from 3 different test images. The performance of Otsu, Chan-Vese, and 3DVascNet was evaluated using these small 3D masks computing both the DC and clDice in 3D.

Statistical Analysis

The Mann-Whitney *U* test was used to compare distributions of vascular morphometrics. More specifically, we used this statistical test to compare the distributions obtained with the manual ground-truth masks and 2D MIP of 3DVascNet masks. The test was applied separately for each vascular morphometric and each intervention group. Furthermore, a Pearson correlation analysis was performed between the values of vascular morphometrics based on the manual ground-truth masks and those based on the 2D MIP of 3DVascNet masks. These 2 analyses were performed to study whether the quantification of vascular parameters obtained from 3DVascNet masks was comparable with those obtained from manual masks.

Moreover, for each vascular morphometric, we used the Mann-Whitney *U* test to compare the distributions of PBS and VEGF (vascular endothelial growth factor), PBS and captopril, PBS and Ang II (angiotensin II), PBS and FLIJ, PBS and FLIO, and PBS and YES. FLIO and FLIJ are code names for datasets of mouse mutants not yet published. This was done separately for the vascular morphometrics computed from 3DVascNet masks and 2D MIP of 3DVascNet masks, to study the differences between 2D and 3D features. All statistical analyses were performed in Python. A *P* value <0.05 was considered statistically significant.

Computer Specifications

The experiments were performed in Python 3.10 on a computer equipped with NVIDIA graphics processing unit GeForce 1080 Ti. The 3D CycleGAN was trained using the open-source DL library Keras.[34] We run all experiments using the source code, and we provide a GUI for researchers without a solid programming background. Microscopy images and segmentation masks shown in the Figures and Supplemental Videos were visualized/analyzed in 3D using Bitplane Imaris.

3DVascNet GUI

A detailed section explaining the installation and step-by-step use of 3DVascNet GUI can be found in the Supplemental Material. Briefly, the directory containing the 3D images must be selected by clicking on the Load the Image(s) button (Figures S4A and S4B). These images will be segmented, and afterward, quantification will be performed upon clicking on the Segment and then Quantify buttons, respectively. During this analysis, a results folder will be automatically created in the images folder (Figure S4C). The segmentation masks and skeletons of the images will be saved in the results folder. These can be visualized using the horizontal and vertical sliders (Figure S4D). Although our software automatically computes the ROI, there is an option to manually select ≥1 ROIs (Select ROI button; Figure S4E). After processing all images, a comma-separated values (CSV) file containing all morphological features for each image will also be saved in the results folder upon clicking on the Save the Results button (Figure S4F). Afterward, the distributions of the features for different groups can be visualized by specifying the group names (separated by commas and without spaces, for instance: captopril,VEGF,PBS), by selecting the features to be visualized, and by using the Visualize Boxplots button (Figure S4G). To load a previously generated CSV file, the Load CSV button can be used at any time, and vasculature features' distributions can be visualized without requiring to perform segmentation and quantification. Finally, the 3D Visualization button allows visualization in 3D of the images or segmentation masks based on the napari viewer.[35]

RESULTS

DL-Based Segmentation Model

Most DL-based approaches require large, manually annotated, ground-truth data sets for training and validation. These annotations should be sufficiently accurate for the model to perform correctly. However, manual segmentation of blood vessels in 3D is demanding. The generation of 3D data sets of sufficient diversity and complexity would require a large effort from researchers. Thus, to overcome the necessity of manually annotating 3D segmentation masks for training, we explored a specific DL method that only requires examples of images from the source domain and target domain for training: the so-called 3D CycleGAN model.[17] We trained the 3D CycleGAN model with 3D microscopy images of retinal blood vessels (the source domain) and synthetic 3D masks (the target domain). The synthetic 3D masks were created based on PolNet,[18] using manual 2D masks of retinal vessels, which were previously generated and published by Barbacena et al.[6] This way, we overcome the burden and need of generating manually annotated 3D masks.

To select the best epoch of the 3D CycleGAN model, we used the NMI metric.[33] Inspired by the work of Zhang et al.,[32] we evaluated the performance of the NMI metric computed between 2D images of blood vessels and 2D masks (ground-truth masks and masks with some perturbations). For each ground-truth segmentation mask, we generated 4 masks by applying dilation and erosion (considering an ellipsoid structuring element of size 5×5), and salt and pepper noise with probability of 0.1 and 0.04. Thereafter, for each image, we computed 5 values of NMI (Equation 9) between the image and each of the 5 masks (the ground-truth mask and the 4 transformed masks). Finally, we computed the accuracy of the NMI metric as the number of times this metric gives a better score (higher value) to the ideal mask (annotated by experts) instead of the transformed mask. Our results showed that NMI is a good unsupervised evaluation metric because it presents high accuracy values, which means that it performs well in giving a better score to the best segmentation mask (Table 2). Remarkably, the accuracy values obtained for NMI in our data set are considerably higher than the ones obtained with the metrics and data sets considered in the study by Zhang et al.[32]

Next, we studied the performance of the 3D CycleGAN model using preprocessing and postprocessing steps. Figure 2A shows the segmentation masks obtained with the 3D CycleGAN model for a raw and a preprocessed image patch (2 masks on the top) and the corresponding postprocessed masks (2 masks on the bottom). These results show that the postprocessing step improves the quality of both segmentation masks by removing noise. However, even after applying the postprocessing step, the segmentation mask obtained with the raw image presents some errors, such as oversegmentation of vessels, indicated with a white arrow. Moreover, our quantitative results on the test images from data set 1 showed that the 3D CycleGAN model tested on the raw and preprocessed images, with the postprocessing step, achieves DC values of 0.8720±0.0367 and 0.9319±0.0288, respectively. Hence, the qualitative and quantitative results suggest that the best segmentation masks are obtained considering both the preprocessing and postprocessing steps. These results show that we generated a DL-based approach that is capable of accurately segmenting blood vessels in 3D without needing 3D manual ground truth for training. We named our approach 3DVascNet. A full-scale image of the segmentation obtained using 3DVascNet is shown in Figure 2B.

Next, we evaluated our method and pipeline by comparing the performance of 3DVascNet with 2 classical methods for image segmentation. Because our data set does not contain 3D ground-truth training masks, we are unable to compare the performance of 3DVascNet with supervised segmentation methods. Thus, we considered 2 unsupervised methods for segmentation: Otsu thresholding[13,26,27] and Chan-Vese.[28] These approaches were applied on the preprocessed images, and the output segmentation masks were subjected to the postprocessing step. For qualitative comparison, we show in Figure 2C an example of a patch extracted from a 3D preprocessed image and corresponding segmentation masks obtained with Otsu thresholding, Chan-Vese, and 3DVascNet. On the one hand, Otsu thresholding is more sensitive to heterogeneous staining resulting in errors in the segmentation mask as shown with white arrows (Figure 2C). On the other hand, Chan-Vese is more susceptible to

Table 2. Accuracy (%) of the Unsupervised Segmentation Evaluation Metric: Normalized Mutual Information (Equation 9)

2D mask with perturbations	Dilation	Erosion	Salt and pepper noise (P=0.1)	Salt and pepper noise (P=0.04)
Original image	93.88%	85.71%	100%	100%
Normalized image	93.88%	100%	100%	100%

2D indicates 2 dimensional.

background noise compared with Otsu thresholding and 3DVascNet (white arrows in Figure 2C). More examples of segmentation masks obtained with the compared approaches are shown in Figure S5. The qualitative results indicate that the best masks are obtained with 3DVascNet (Figure 2C; Figure S5). The quantitative results showed that 3DVascNet achieves the highest DC and clDice values and outperforms traditional segmentation methods (Table 3).

To further validate the performance of 3DVascNet, we created ground-truth masks in 3D. Highlighting the burden of manual creation of 3D ground truth, it took ≈40 hours for the creation of these 9 small masks, whereas 3DVascNet segmented the corresponding image patches in about 3 minutes. The performance of Otsu, Chan-Vese, and 3DVascNet was evaluated in 3D using these masks. The results indicate that the best DC and clDice values are obtained with the 3DVascNet method (Table 4), which confirms the robustness of 3DVascNet to segment vessels in 3D and its superiority compared with traditional segmentation approaches. Moreover, the high clDice value obtained with 3DVascNet, ≈92%, is a strong indicator of the great similarity between the topology of the 3D ground-truth masks and 3DVascNet masks, which is also consistent with the qualitative results shown in Figure 3. Thus, we have developed an efficient method for training and segmenting blood vessels in 3D without the need for manual labor to create ground-truth information.

Generalization Capability

We next evaluated the generalization capability of 3DVascNet. First, we measured the performance of 3DVascNet on additional data sets of retinal vascular networks, acquired with different vessel markers, microscopes, and by different users (data sets 2–4 described in Table S1). Figure 4A represents the segmentation masks obtained with 3DVascNet for 3 patches taken from images from data sets 2 to 4. These results show that 3DVascNet provides segmentation masks with high quality for images acquired under different conditions and on different microscopes. This is confirmed by comparing 3DVascNet with Otsu and Chan-Vese methods (Table 5). In fact, 3DVascNet outperforms the traditional methods for segmentation, presenting the highest DC mean and smallest SD for all data sets considered. The 3DVascNet model achieves mean DC and clDice values of ≈93% and ≈96% for these 3 data sets, while Otsu thresholding achieves mean DC and clDice values of ≈88% and ≈91%, and Chan-Vese achieves only ≈76% for both metrics. Hence, 3DVascNet presents a high generalization capability.

Next, we qualitatively studied the generalization capability of the proposed segmentation model using additional experimental conditions. First, we tested whether 3DVascNet could perform well using CD31 (cluster of differentiation 31) staining instead of ICAM2 (intercellular adhesion molecule 2), which was used for the previous data sets (1–4). Indeed, CD31 staining was also well segmented as ICAM2, both at the same stage of retinal development (postnatal day 6) as the previous data sets (Figure 4B). 3DVascNet also segments well 3 layers of blood vessels corresponding to a later stage (postnatal day 20), where the 3 dimensionality of the network is more apparent (Figure 4B; Videos S2 through S4). In addition, we segmented a thick section from kidney medulla stained for CD31. Remarkably, 3DVascNet also performed robustly and correctly segmented almost all individual vessels. Finally, we tested 3DVascNet on various adult retinas that exhibited differences in staining intensity/quality across different layers (Videos S5 through S13; Figure 4C). Videos S7, S10, and S13 show that 3DVascNet also demonstrates excellent 3D vasculature segmentation in these adult retinas.

Thus, 3DVascNet presents a good ability to segment vessels, even in the presence of image variations resulting from differences in staining techniques, microscopy equipment, image modality, and mouse organ and its developmental stage. Training with further data sets will likely further expand its generalization capabilities.

To facilitate the usage of 3DVascNet, we have generated a GUI that has been developed in Python (Figure S4). 3DVascNet's GUI allows automatic analysis of 3D microscopy images of blood vessels. 3DVascNet has been trained only with images from retinal blood vessels at postnatal day 6 of development. Thus, it should perform well on similar data sets. Yet, other vascular networks can be tested using our software.

Benchmarking 3DVascNet for Vascular Morphometrics

Next, we tested whether the quantification of several vascular parameters (vascular morphometrics), such as vessel density, vessel radius, number of branching points, and vessel length, using 3DVascNet was comparable with vascular morphometrics obtained from manual segmentation. To be able to have similar dimensions and units, we calculated vascular morphometrics from 2D MIP of the 3DVascNet masks and compared them with the vascular morphometrics computed from the 2D ground-truth masks. Remarkably, the distributions for the different vascular morphometrics obtained through 3DVascNet are similar to the distributions obtained using user-dependent

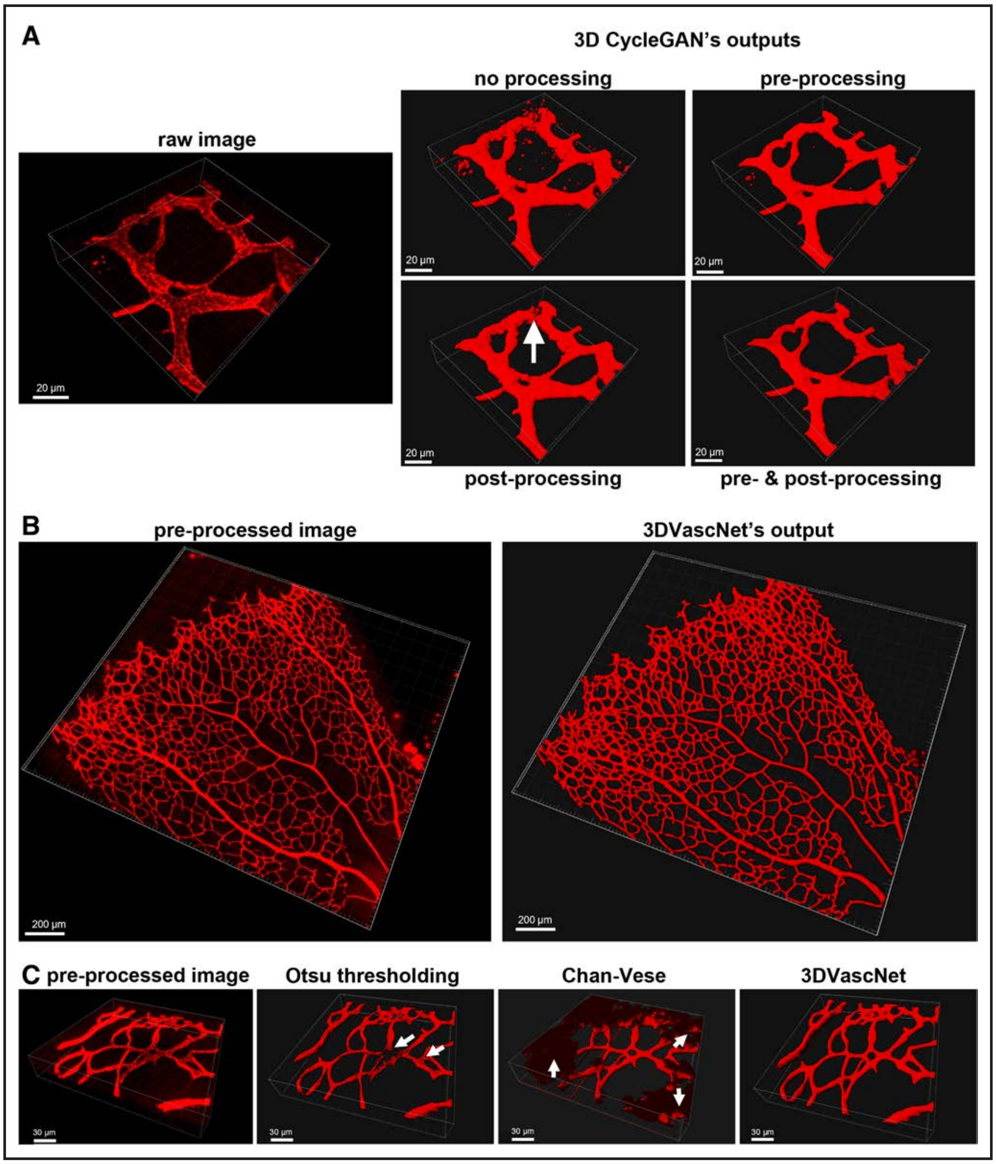

Figure 2. Optimization of the proposed segmentation model and comparative performance.
A, Demonstration of the effects of preprocessing and postprocessing steps in the output of 3-dimensional (3D) CycleGAN using as an example a 3D patch extracted from the input image (**left**) and the corresponding segmentation masks obtained with the 3D CycleGAN (**right**), corresponding to the different processing steps. **B**, Example of the segmentation of an entire postnatal day 6 retinal vascular network, labeled by ICAM2 (intercellular adhesion molecule 2), using 3DVascNet. Preprocessed image (**left**), and 3DVascNet segmentation mask (**right**). **C**, Examples of a 3D patch extracted from a preprocessed image (**left**) and the corresponding segmentation masks obtained with Otsu thresholding, Chan-Vese, and 3DVascNet.

2D ground truth (Figure 5). Importantly, there are no statistical differences ($P>0.05$), which confirms that results are equivalent between both approaches. To further confirm that results are within the same range, we also performed a Pearson correlation analysis between the values of vascular morphometrics based on the manual ground-truth masks and those based on the 3DVascNet masks. We obtained high R^2 values for all features ($R^2 \approx 0.9$), which demonstrates a linear relationship between 2 variables.

These findings demonstrate that 3DVascNet is a powerful automated method for blood vessel segmentation and quantification yielding robust scientific outcomes that are within the range of manual segmentation.

3D Quantification of Vascular Morphometrics

Next, we tested whether we could obtain additional information from our data sets when using quantifications in 3D instead of quantifications from 2D MIP images. To do so, we computed 3D skeletons using the method proposed in the study by Lee et al[24] (Figure 6A). We used the 3D skeletons to calculate the 4 previously characterized vascular morphometric parameters in 3D

Table 3. DC and clDice Values of the Compared Vessel Segmentation Approaches on Data Set 1, Considering the Preprocessing and Postprocessing Steps

Method		Data set 1	DC (mean±SD)	clDice (mean±SD)
Traditional segmentation methods	Otsu thresholding	Train	0.9095±0.0341	0.9426±0.0253
		Test	0.8750±0.0539	0.8812±0.0764
	Chan-Vese	Train	0.8484±0.0787	0.8674±0.0528
		Test	0.8014±0.1416	0.8039±0.1751
3DVascNet		Train	0.9450±0.0240*	0.9788±0.0190*
		Test	0.9319±0.0288*	0.9549±0.0402*

All values correspond to mean±SD. clDice indicates centerline Dice; and DC, Dice coefficient.

*The best results.

(Figure 6B). Interestingly, our results demonstrate that some quantifications, and some relative differences between experimental groups, differ when quantifying parameters in 2D or 3D, while in other cases, tendencies are maintained (Figure 6B). For instance, the vessel density in 3D follows the same trends as 2D quantifications. However, we must note that absolute values differ much, and 2D quantifications tend to suggest a much higher vessel density than in reality.

Importantly, the mean vessel radius was substantially changed. Notably, higher mean vessel radius values were observed in 2D compared with 3D. Also, in 3D, most of the conditions do not present a significant difference, regarding mean vessel radius, compared with PBS, while in 2D, there are significant differences ($P<0.05$) between PBS and captopril, PBS and FLIO, and PBS and Yes (Figure 6B). We attributed this difference to the elliptical cross-section of retinal vessels. The long axis of the vessels is parallel to the inner limiting membrane, that is, parallel with respect to the XY plane (Figure S6, top). This leads to an overestimation of vessel radius when using 2D MIPs (Figure S6, bottom) and better estimation of vessel radius in 3D because in 3D, we quantify the mean radius across the entire contour of the vessel cross-section.

Mean branch length and branching point density followed similar trends to 2D quantifications, yet some of the relative changes, when compared with the PBS group, did not correlate well between 2D and 3D quantifications. For instance, statistically significant differences ($P<0.05$) in the mean branch length between PBS and Ang II and PBS and FLIO are not evident in 3D. Furthermore, our 3D analysis indicated statistically significant differences ($P<0.05$) in branching point density between PBS and captopril and PBS and Ang II, and in the mean branch length between PBS and VEGF, which were not evident in the 2D study (Figure 6B).

Table 4. DC and clDice Values of the Compared Vessel Segmentation Approaches for the 9 Image Patches for Which 3D Ground-Truth Masks Were Manually Generated for Evaluation

	3D DC	3D clDice
Otsu	0.7014±0.0958	0.8535±0.0321
Chan-Vese	0.5301±0.2315	0.5682±0.2937
3DVascNet	0.7344±0.0620*	0.9173±0.0389*

Here, the images after applying the percentile-based normalization were considered. 3D indicates 3 dimensional; clDice, centerline Dice; and DC, Dice coefficient.

*The best results.

Thus, we demonstrate that 3D quantifications are important to accurately describe vascular morphometrics and that caution must be taken when drawing conclusions from 2D MIPs.

DISCUSSION

Analysis of 3D vascular networks is an essential step to understand several physiological and pathological processes. In this work, we disclose 3DVascNet, a novel DL-based software for automated segmentation and quantification of 3D retinal vascular networks. Our approach is based on the 3D CycleGAN model that translates 3D images into 3D segmentation masks. Importantly, the proposed method for segmentation only required 2D ground-truth masks to train the 3DVascNet model. The main advantage of the proposed pipeline compared with previous approaches[4–9] is that our method can be used to study retinas at any developmental stage because it allows to perform segmentation in 3D. During the later phases (postnatal day 8 and beyond), the blood vessels in the retina develop a 3-layered 3D structure. Analyzing in 2D retinas from these stages can lead to incorrect estimation of branching points and branch lengths due to the overlap of vessels from different layers. We have shown that our approach is able to segment 3D blood vessels corresponding to a postnatal day 20 retina (Videos S2 through S4) and adult retinas (Videos S5 through S13). Thus, the proposed method will clearly be a valuable tool in vascular research.

Even if 3DVascNet produces fully automated high-quality 3D segmentation masks, we strongly recommend the acquisition of microscopy images without the presence of small noisy regions of high intensity, typically caused by the presence of dirt or air bubbles, as these were incorrectly segmented by our model, compromising the quality of the segmentation masks (Figure S7).

We have shown that our method outperforms traditional approaches for segmentation (Otsu thresholding[26] and Chan-Vese[28]). On the one hand, these classical methods are based on image intensities, resulting in oversegmentation in regions with background noise and undersegmentation in regions with heterogeneous staining. On the other hand, the 3DVascNet performs segmentation based on automatically extracted features from the input images, which are learned throughout the training process. Our results show that DL-based segmentation

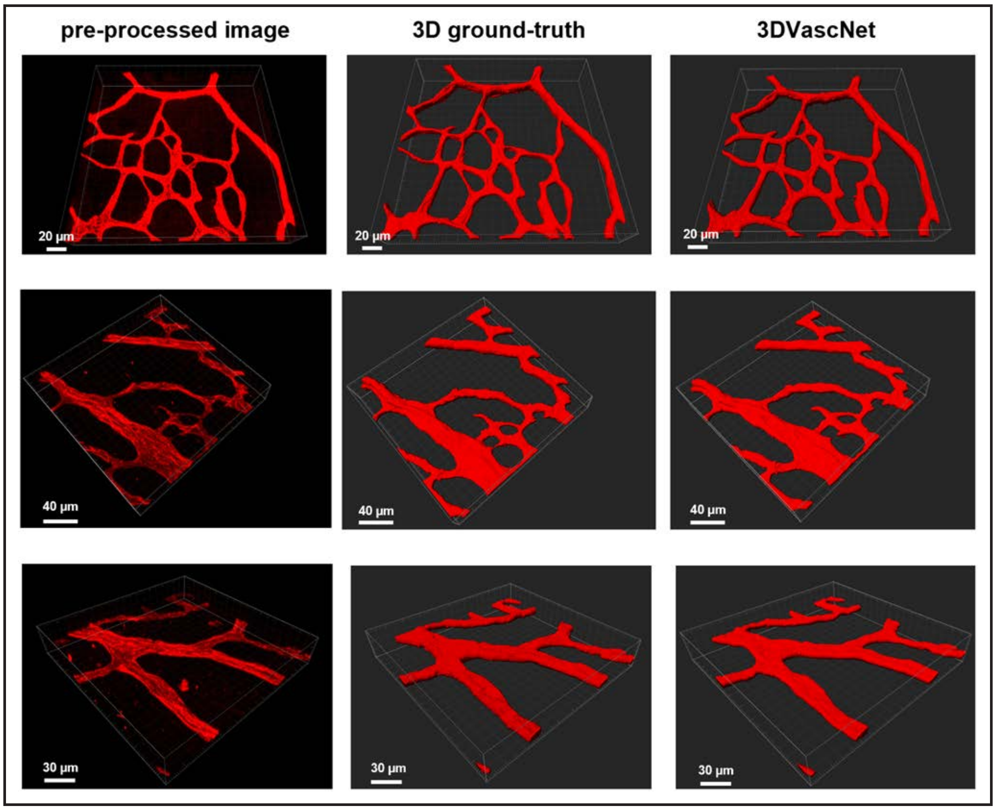

Figure 3. Comparison of 3-dimensional (3D) ground-truth and 3DVascNet masks.
Preprocessed image patches (**left**), corresponding manually created 3D ground-truth masks (**middle**), and 3DVascNet masks (**right**).

masks represent accurately the vessels in the microscopy images. Furthermore, our results indicate that the proposed approach presents a high generalization capability. Hence, it can be used for accurate segmentation of 3D retinal vessels in microscopy images acquired under various different experimental conditions.

Compared with the method proposed in the study by Bernabeu et al,[18] the proposed approach produces higher quality 3D segmentation masks. This is likely because in this work, the 3D segmentation masks are obtained from the 3D images, whereas in the study by Bernabeu et al,[18] the 3D masks are computed based on 2D masks and a set of assumptions, such as the spherical shape of the vessel cross-section and coplanarity of the network. Although 3DVascNet presented a high generalization capability, it may be necessary to further improve its performance by retraining with new images. Since the source code of 3DVascNet and the data sets are made publicly available, our model can be easily trained with user-specific 3D images and provided 3D masks, which will extend 3DVascNet to other vascular networks. For instance, light-sheet microscopy has been used to image entire vascular networks of the retina[36,37] and the brain.[38] It would be interesting to test and expand 3DVascNet to segment these types of data sets.

Although the source codes of most of the described works[6,8–10] are made publicly available, only 2 works[7,8] have presented a GUI, which can easily be used by researchers without a solid expertise in computer science and programming. To overcome this issue, we present not only 3DVascNet source code, which can be easily adapted for the segmentation of other vascular networks, but also a user-friendly graphical interface. However, the main limitation is that 3DVascNet must be executed on a powerful computer with an NVIDIA graphics processing unit.

We have validated the performance of 3DVascNet by comparing its output to the quantification of 4 vascular morphometric parameters based on ground-truth segmentation masks. Our results indicated that there are no statistically significant differences between the vascular morphometric parameters obtained from the manual masks and 2D projections of 3DVascNet masks. Also, there was a high linear correlation between manual and 3DVascNet-based computation of vascular morphometric parameters ($R^2 \geq 0.87$). Thus, the quantitative analyses yielded similar results compared with manual segmentation. However, 2D analysis of retinal vasculature has several limitations. For instance, 2D studies may be misleading due to false-positive branches and branching points resulting from crossings of various vessels along the *z* axis. This is due to the fact that the 2D quantifications do not take into consideration the true geometry of the vessels. Hence, we also computed the

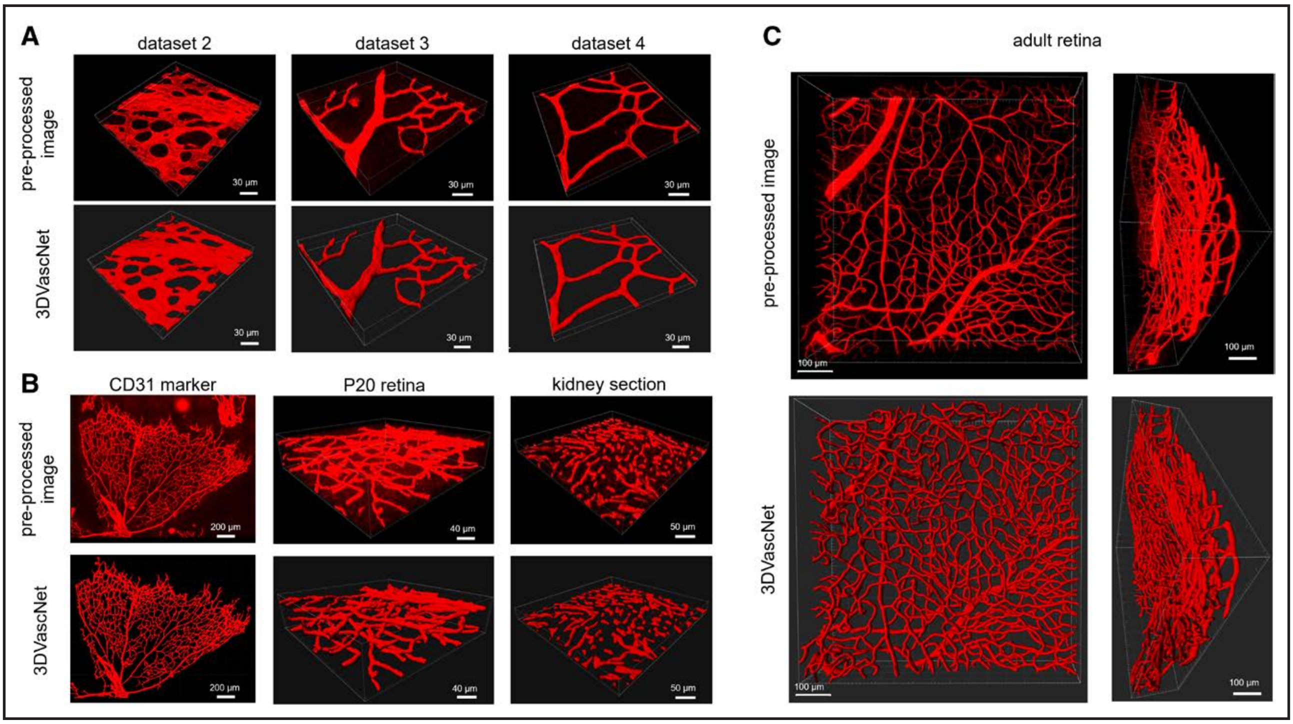

Figure 4. Generalization capability of 3DVascNet.
A, Examples of 3-dimensional (3D) patches extracted from the preprocessed images (**top**), and the corresponding segmentation masks obtained with 3DVascNet (**bottom**), corresponding to data sets 2, 3, and 4 (from left to right, respectively). **B**, Examples of 3D patches extracted from the preprocessed images (**top**) and the corresponding segmentation masks obtained with 3DVascNet (**bottom**), corresponding to different conditions (CD31 [cluster of differentiation 31] marker and postnatal day 20 [P20] retina), and a different organ (kidney section). **C**, Example of an adult retina segmentation: preprocessed image in 2 different perspectives (**top**), and corresponding mask obtained with 3DVascNet also in the same 2 perspectives (**bottom**).

3D features from the 3DVascNet masks and corresponding skeletons. Although the ranges of branch lengths across different intervention groups are similar in 3D and 2D, there is a notable disparity in the quantification of the vessel radius between 2D and 3D. These differences are likely explained by the assumption that vessel cross-sections are circular, while, in reality, they are oval, leading to an overestimation of vessel radii in 2D. Thus, the 3D quantifications obtained with the proposed approach are superior in comparison to the traditional 2D analysis of retinal blood vessels. This is due to their ability to capture the volumetric information providing a more realistic quantification of the 3D blood vessels. To conclude, the differences between the 2D and 3D quantifications enhance the need of undertaking 3D analysis of the vasculature instead of studying its 2D projections.

To summarize, in this work, we have presented 3DVascNet, a software for segmentation and quantification of blood vessels in 3D. 3DVascNet is fully automated; thus it will be used to perform large-scale analysis in 3D, which will be of great importance to answer key open questions in vascular biology.

Table 5. DC and clDice Values of the Compared Vessel Segmentation Approaches on Data Sets 2, 3, and 4, Considering the Images After Applying the Percentile-Based Normalization

		Data set 2	Data set 3	Data set 4	All
Otsu	DC	0.9260±0.0280	0.8306±0.1573	0.9050±0.0172	0.8798±0.1111
	clDice	0.9527±0.0249	0.8755±0.1658	0.9134±0.0374	0.9098±0.1133
Chan-Vese	DC	0.7947±0.1271	0.6617±0.2494	0.8738±0.0547	0.7575±0.1971
	clDice	0.7727±0.1523	0.6643±0.2709	0.8906±0.0657	0.7557±0.2151
3DVascNet	DC	0.9486±0.0186*	0.9230±0.0816*	0.9380±0.0109*	0.9349±0.0545*
	clDice	0.9778±0.0124*	0.9513±0.0864*	0.9599±0.0200*	0.9620±0.0576*

All values correspond to mean±SD. clDice indicates centerline Dice; and DC, Dice coefficient.
*The best results.

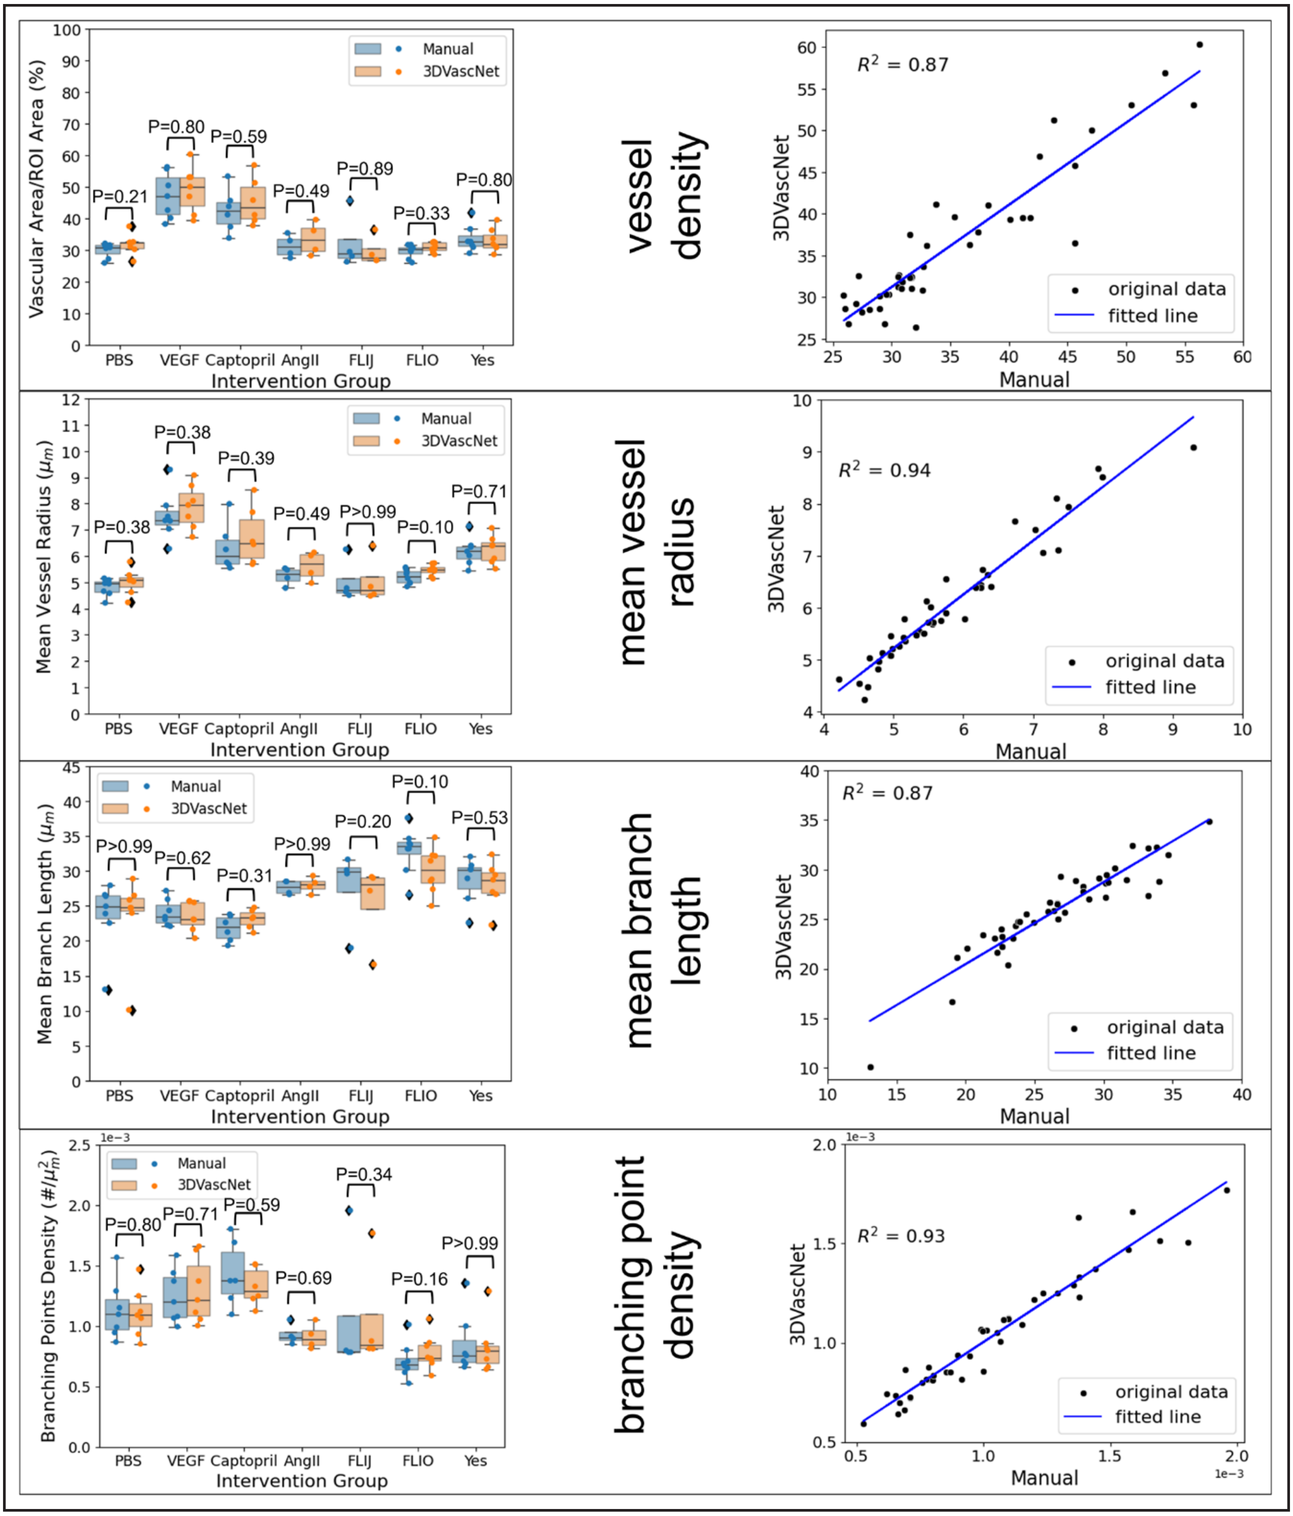

Figure 5. **Performance of 3DVascNet using vascular morphometrics.**

Box plots (left graph) represent side-by-side comparison of vascular morphometrics computed from 2-dimensional maximum-intensity projection of manual segmentation masks (in blue) or 3DVascNet masks (in orange). Statistical analysis between the distributions of manual and automated features was performed, for each vascular morphometric and each intervention group, using the Mann-Whitney *U* test; the corresponding *P* values are presented above each pair of distributions compared. A *P* value <0.05 was considered statistically significant. Pearson correlation analysis (right graph) between the vascular morphometrics based on manual masks (*x* axis) and those based on the 3DVascNet masks (*y* axis). Data represent vessel density (vascular area/region of interest [ROI] area [%]), mean vessel radius (μm), mean branch length (μm), and branching point density (#/μm^2). For correlation analyses, we considered 43 images, corresponding to the groups shown in the box plots. Here, we considered only the intervention groups that have at least 4 images: PBS (n=7), VEGF (vascular endothelial growth factor; n=7), captopril (n=6), Ang II (angiotensin II; n=4), FLIJ (n=4), FLIO (n=8), and Yes (n=7).

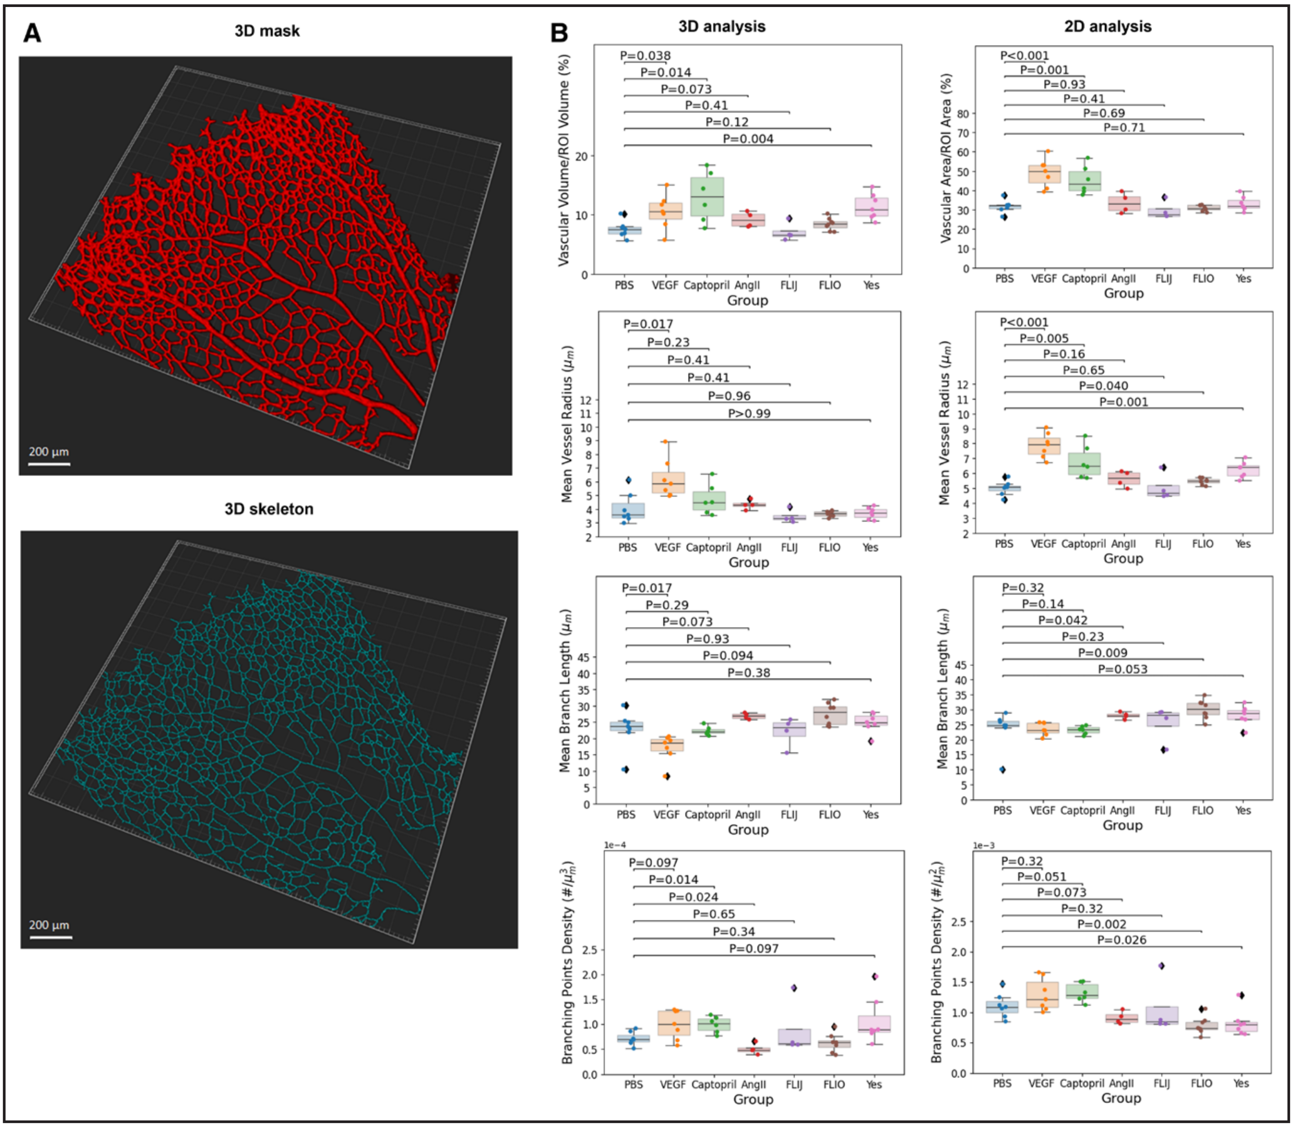

Figure 6. Comparison between 2-dimensional (2D) and 3-dimensional (3D) vascular morphometrics.
A, Representative 3D segmentation mask from an Ang II (angiotensin II) retina (**top**) and corresponding 3D skeleton (**bottom**) computed based on the segmentation mask. **B**, Boxplots represent vascular morphometrics computed from 3DVascNet masks (left graph) or 2D maximum-intensity projection of 3DVascNet masks (right graph), corresponding to the 4 studied morphometrics: vessel density (%), mean vessel radius (μm), mean branch length (μm), and branching point density (#/μm^3 or #/μm^2). Here, we also considered only the intervention groups that have at least 4 images: PBS (n=7), VEGF (vascular endothelial growth factor; n=7), captopril (n=6), Ang II (n=4), FLIJ (n=4), FLIO (n=8), and Yes (n=7). Statistical analysis between the distributions (PBS, VEGF), (PBS, captopril), (PBS, Ang II), (PBS, FLIJ), (PBS, FLIO), and (PBS, YES) was performed, for each vascular morphometric, using the Mann-Whitney *U* test; the corresponding *P* values are shown above each pair of distributions compared. A *P* value <0.05 was considered statistically significant.

ARTICLE INFORMATION

Received February 19, 2024; accepted May 7, 2024.

Affiliations

Instituto de Sistemas e Robótica, LARSyS, Instituto Superior Técnico (H.N., M.S.) and Instituto de Medicina Molecular-João Lobo Antunes, Faculdade de Medicina (H.N., C.A.F.), Universidade de Lisboa, Lisbon, Portugal. Universidade Católica Portuguesa, Católica Medical School, Católica Biomedical Research Centre, Lisbon, Portugal (H.N., C.A.F.).

Acknowledgments

The authors thank Rui Benedito (Centro Nacional de Investigaciones Cardiovasculares Carlos III) and Lena Claesson-Welsh and Yi Jin (Uppsala University) for providing images or retinas used in this study. They thank Miguel Bernabeu (Edinburgh University) for insights and comments on the usage of PolNet. They are grateful to Ana Figueiredo, Andreia Pena, Daniela Ramalho, and Marta Saraiva for their valuable suggestions on the graphical user interface. They thank Pedro Barbacena, Marie Ouarné, Ana Figueiredo, Andreia Pena, and Daniela Ramalho for providing the microscopy images used in this study. They thank Pedro Barbacena, Andreia Pena, and Daniela Ramalho for manually annotating the 2-dimensional ground-truth masks used in this work. H. Narotamo had full access to all the data in the study and takes responsibility for its integrity and the data analysis.

Sources of Funding

H. Narotamo was supported by Fundação para a Ciência e a Tecnologia (FCT) Doctoral Grant 2020.04511.BD (doi: 10.54499/2020.04511.BD). H. Narotamo and M. Silveira were supported by LARSyS funding (doi: 10.54499/LA/P/0083/2020, 10.54499/UIDP/50009/2020, and 10.54499/UIDB/50009/2020). C.A. Franco was supported by the European Research

Council starting grant (679368), LaCaixa Foundation grant (ECARD, HR22-00551), a grant from Fondation LeDucq (17CVD03), and the FCT funding (grants PTDC/MEDPAT/31639/2017, PTDC/BIA-CEL/32180/2017, and CEECIND/02589/2018).

Disclosures

None.

Supplemental Material

Figures S1–S7
Table S1
Supporting Information
Videos S1–S13
Major Resources Table
Reference 39

REFERENCES

1. Potente M, Gerhardt H, Carmeliet P. Basic and therapeutic aspects of angiogenesis. *Cell*. 2011;146:873–887. doi: 10.1016/j.cell.2011.08.039
2. Potente M, Mäkinen T. Vascular heterogeneity and specialization in development and disease. *Nat Rev Mol Cell Biol*. 2017;18:477–494. doi: 10.1038/nrm.2017.36
3. Meng X, Xing Y, Li J, Deng C, Li Y, Ren X, Zhang D. Rebuilding the vascular network: in vivo and in vitro approaches. *Front Cell Dev Biol*. 2021;9:639299. doi: 10.3389/fcell.2021.639299
4. Milde F, Lauw S, Koumoutsakos P, Iruela-Arispe ML. The mouse retina in 3D: quantification of vascular growth and remodeling. *Integr Biol (Camb)*. 2013;5:1426–1438. doi: 10.1039/c3ib40085a
5. Rust R, Grönnert L, Dogançay B, Schwab ME. A revised view on growth and remodeling in the retinal vasculature. *Sci Rep*. 2019;9:1–9. doi: 10.1038/s41598-019-40135-2
6. Barbacena P, Dominguez-Cejudo M, Fonseca CG, Gómez-González M, Faure LM, Zarkada G, Pena A, Pezzarossa A, Ramalho D, Giarratano Y, et al. Competition for endothelial cell polarity drives vascular morphogenesis in the mouse retina. *Dev Cell*. 2022;57:2321–2333.e9. doi: 10.1016/j.devcel.2022.09.002
7. Zudaire E, Gambardella L, Kurcz C, Vermeren S. A computational tool for quantitative analysis of vascular networks. *PLoS One*. 2011;6:e27385. doi: 10.1371/journal.pone.0027385
8. Montoya-Zegarra JA, Russo E, Runge P, Jadhav M, Willrodt AH, Stoma S, Nørrelykke SF, Detmar M, Halin CA. A novel software for the automated morphometric analysis of vascular networks in tissues. *Angiogenesis*. 2019;22:223–236. doi: 10.1007/s10456-018-9652-3
9. Wang X, Zhu G, Wang S, Rhen J, Pang J, Zhang Z. Vessel tech: a high-accuracy pipeline for comprehensive mouse retinal vasculature characterization. *Angiogenesis*. 2021;24:7–11. doi: 10.1007/s10456-020-09752-8
10. Todorov MI, Paetzold JC, Schoppe O, Tetteh G, Shit S, Efremov V, Todorov-Völgyi K, Düring M, Dichgans M, Piraud M, et al. Machine learning analysis of whole mouse brain vasculature. *Nat Methods*. 2020;17:442–449. doi: 10.1038/s41592-020-0792-1
11. Chen C, Chuah JH, Ali R, Wang Y. Retinal vessel segmentation using deep learning: a review. *IEEE Access*. 2021;9:111985–112004. doi: 10.1109/access.2021.3102176
12. Fraz MM, Remagnino P, Hoppe A, Uyyanonvara B, Rudnicka AR, Owen CG, Barman SA. Blood vessel segmentation methodologies in retinal images: a survey. *Comput Methods Programs Biomed*. 2012;108:407–433. doi: 10.1016/j.cmpb.2012.03.009
13. Khandouzi A, Ariafar A, Mashayekhpour Z, Pazira M, Baleghi Y. Retinal vessel segmentation, a review of classic and deep methods. *Ann Biomed Eng*. 2022;50:1292–1314. doi: 10.1007/s10439-022-03058-0
14. Stahl A, Connor KM, Sapieha P, Chen J, Dennison RJ, Krah NM, Seaward MR, Willett KL, Aderman CM, Guerin KI, et al. The mouse retina as an angiogenesis model. *Invest Ophthalmol Vis Sci*. 2010;51:2813–2826. doi: 10.1167/iovs.10-5176
15. Srinidhi CL, Aparna P, Rajan J. Recent advancements in retinal vessel segmentation. *J Med Syst*. 2017;41:1–22. doi: 10.1007/s10916-017-0719-2
16. Fruttiger M. Development of the retinal vasculature. *Angiogenesis*. 2007;10:77–88. doi: 10.1007/s10456-007-9065-1
17. Zhu JY, Park T, Isola P, Efros AA. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision. 2017:2223–2232.
18. Bernabeu MO, Jones ML, Nash RW, Pezzarossa A, Coveney PV, Gerhardt H, Franco CA. PolNet: a tool to quantify network-level cell polarity and blood flow in vascular remodeling. *Biophys J*. 2018;114:2052–2058. doi: 10.1016/j.bpj.2018.03.032
19. Kumar KS, Singh NP. Analysis of retinal blood vessel segmentation techniques: a systematic survey. *Multimedia Tools Appl*. 2022;82:7679–7733. doi: 10.1007/s11042-022-13388-9
20. Isola P, Zhu JY, Zhou T, Efros AA. Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2017:1125–1134.
21. Brownlee J. How to develop a CycleGAN for image-to-image translation with Keras. 2019. https://machinelearningmastery.com/cyclegan-tutorial-with-keras/
22. Kingma DP, Ba J. Adam: a method for stochastic optimization. In: Proceedings of the 3rd International Conference on Learning Representations. 2015.
23. Moreira A, Santos MY. Concave hull: a k-nearest neighbours approach for the computation of the region occupied by a set of points. In: Proceedings of the International Conference on Computer Graphics Theory and Applications. INSTICC Press; 2007:61–68.
24. Lee TC, Kashyap RL, Chu CN. Building skeleton models via 3-D medial surface axis thinning algorithms. *CVGIP Graph Models Image Process*. 1994;56:462–478. doi: 10.1006/cgip.1994.1042
25. Nunez-Iglesias J, Blanch AJ, Looker O, Dixon MW, Tilley L. A new Python library to analyse skeleton images confirms malaria parasite remodelling of the red blood cell membrane skeleton. *PeerJ*. 2018;6:e4312. doi: 10.7717/peerj.4312
26. Otsu N. A threshold selection method from gray-level histograms. *IEEE Trans Syst Man Cybern*. 1979;9:62–66. doi: 10.1109/tsmc.1979.4310076
27. Yousefi J. *Image Binarization Using Otsu Thresholding Algorithm*. University of Guelph. 2011.
28. Chan TF, Vese LA. Active contours without edges. *IEEE Trans Image Process*. 2001;10:266–277. doi: 10.1109/83.902291
29. Dice LR. Measures of the amount of ecologic association between species. *Ecology*. 1945;26:297–302. doi: 10.2307/1932409
30. Moccia S, Momi ED, Hadji SE, Mattos LS. Blood vessel segmentation algorithms-review of methods, datasets and evaluation metrics. *Comput Methods Programs Biomed*. 2018;158:71–91. doi: 10.1016/j.cmpb.2018.02.001
31. Shit S, Paetzold JC, Sekuboyina A, Ezhov I, Unger A, Zhylka A, Pluim JP, Bauer U, Menze BH. clDice-a novel topology-preserving loss function for tubular structure segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021:16560–16569.
32. Zhang H, Fritts JE, Goldman SA. Image segmentation evaluation: a survey of unsupervised methods. *Comput Vis Image Underst*. 2008;110:260–280. doi: 10.1016/j.cviu.2007.08.003
33. Studholme C, Hill DL, Hawkes DJ. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognit*. 1999;32:71–86. doi: 10.1016/s0031-3203(98)00091-0
34. Chollet F. Keras. 2015. https://github.com/fchollet/keras
35. Chiu CL, Clack N; The Napari Community. Napari: a python multidimensional image viewer platform for the research community. *Microsc Microanal*. 2022;28:1576–1577. doi: 10.1017/s1431927622006328
36. Darche M, Verschueren A, Belle M, Boucherit L, Fouquet S, Sahel JA, Chédotal A, Cascone I, Paques M. Three-dimensional characterization of developing and adult ocular vasculature in mice using in toto clearing. *Commun Biol*. 2022;5:1135. doi: 10.1038/s42003-022-04104-2
37. Prahst C, Ashrafzadeh P, Mead T, Figueiredo A, Chang K, Richardson D, Venkaraman L, Richards M, Russo AM, Harrington K, et al. Mouse retinal cell behaviour in space and time using light sheet fluorescence microscopy. *Elife*. 2020;9:e49779. doi: 10.7554/eLife.49779
38. Kirst C, Skriabine S, Vieites-Prado A, Topilko T, Bertin P, Gerschenfeld G, Verny F, Topilko P, Michalski N, Tessier-Lavigne M, et al. Mapping the fine-scale organization and plasticity of the brain vasculature. *Cell*. 2020;180:780–795.e25. doi: 10.1016/j.cell.2020.01.028
39. Jin Y, Ding Y, Richards M, Kaakinen M, Giese W, Baumann E, Szymborska A, Rosa A, Nordling S, Schimmel L, et al. Tyrosine-protein kinase Yes controls endothelial junctional plasticity and barrier integrity by regulating VE-cadherin phosphorylation and endocytosis. *Nat Cardiovasc Res*. 2022;1:1156–1173. doi: 10.1038/s44161-022-00172-z