

Physiology-based simulation of the retinal vasculature enables annotation-free segmentation of OCT angiographs

Martin J. Menten^{1,2}, Johannes C. Paetzold^{1,2}, Alina Dima¹, Bjoern H. Menze³, Benjamin Knier⁴, and Daniel Rueckert^{1,2}

¹ Lab for AI in Medicine, Klinikum rechts der Isar, Technical University of Munich

² BioMedIA, Department of Computing, Imperial College London

³ Department of Quantitative Biomedicine, University of Zurich

⁴ Department of Neurology, Klinikum rechts der Isar, Technical University of Munich

Abstract. Optical coherence tomography angiography (OCTA) can non-invasively image the eye’s circulatory system. In order to reliably characterize the retinal vasculature, there is a need to automatically extract quantitative metrics from these images. The calculation of such biomarkers requires a precise semantic segmentation of the blood vessels. However, deep-learning-based methods for segmentation mostly rely on supervised training with voxel-level annotations, which are costly to obtain. In this work, we present a pipeline to synthesize large amounts of realistic OCTA images with intrinsically matching ground truth labels; thereby obviating the need for manual annotation of training data. Our proposed method is based on two novel components: 1) a physiology-based simulation that models the various retinal vascular plexuses and 2) a suite of physics-based image augmentations that emulate the OCTA image acquisition process including typical artifacts.

In extensive benchmarking experiments, we demonstrate the utility of our synthetic data by successfully training retinal vessel segmentation algorithms. Encouraged by our method’s competitive quantitative and superior qualitative performance, we believe that it constitutes a versatile tool to advance the quantitative analysis of OCTA images.

1 Introduction

Optical coherence tomography angiography (OCTA) allows the non-invasive acquisition of high-fidelity volumetric images of the eye’s circulatory system [29]. It has become an important tool for ophthalmologists to diagnose and monitor ocular diseases that manifest themselves as pathological changes to the eye’s vessels [9]. Furthermore, the retina has been described as “window to the body” as alterations to its vasculature have been linked to various neurological and cardiac diseases [13,16]. In order to improve the understanding and treatment of these conditions, it is crucial to automatically extract quantitative metrics that characterize the retinal vasculature. The calculation of many of these biomarkers requires a precise semantic segmentation of the blood vessels.

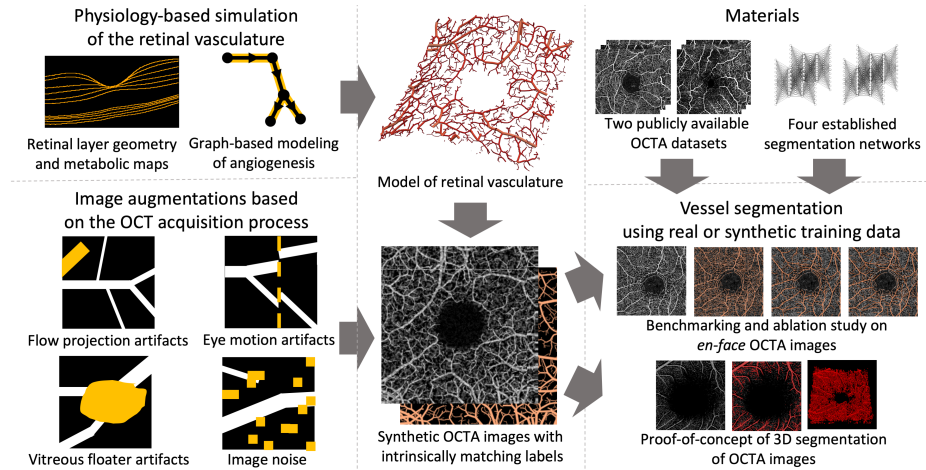


Fig. 1. Overview of the proposed method and conducted experiments.

Recently, deep learning has seen wide-ranging success in the semantic segmentation of medical images, including the segmentation of blood vessels in angiographs. However, training of these algorithms is mostly based on large sets of annotated data, which are time-consuming to curate. Creating voxel-level annotations of the complex branching retinal vasculature in OCTA images is particularly cumbersome. Until now, only two publicly available OCTA datasets with voxel-wise annotations have been released [12,17]. With 226 and 500 images, respectively, these datasets are several orders of magnitude smaller than the largest medical image datasets and a dataset with three-dimensional annotations has not been published yet. There is a clear need for new strategies that are independent of manual labeling to further advance the quantitative analysis of OCTA images.

To address this need, we present a pipeline that can synthesize large amounts of realistic OCTA images with intrinsically matching ground truth labels (see figure 1). Our proposed method is based on two novel components: 1) a physiology-based simulation that models the various retinal vascular plexuses and 2) a suite of physics-based image augmentations that emulate the OCTA image acquisition process including typical artifacts.

We evaluate the utility of our synthetically generated data for the training of supervised vessel segmentation algorithms. In extensive benchmarking experiments involving four neural network architectures and two publicly available OCTA datasets, we compare the segmentation results when training on synthetic OCTA image-label pairs versus real OCTA images with expert-derived annotations. Additionally, we quantify the intrinsic scalability of our approach and investigate how it can facilitate segmentation of the retinal vasculature in three-dimensional OCTA images.

2 Related works

Deep-learning-based segmentation of OCTA images Similar to most segmentation tasks, convolutional neural networks constitute the top-performing methods for segmentation of tubular structures [4]. However, the adoption of supervised deep learning for retinal vessel segmentation in OCTA images has been limited due to a lack of annotated data [11,12,17]. Beyond supervised learning, Liu *et al.* have recently presented an unsupervised method for OCTA segmentation that does not rely on expert-derived labels [15]. However, their method requires multiple scans of the same subject acquired with scanners from different manufacturers, which are usually unobtainable in clinical practice.

Learning with synthetic data One way to overcome data sparsity in machine learning is transfer learning from synthetically generated data. Here, models are pretrained on artificial data and then refined on a small set of real image-label pairs. In medical image segmentation, this approach has been successfully applied to the segmentation of brain tumors [14], endotracheal tubes [3] and catheters [7] as well as tubular structures, such as the whole mouse brain vasculature [32], the dermal vasculature [6] and neurons [20].

Physics-based augmentations Another strategy to mitigate the impact of sparse data is image augmentation, where additional training samples are generated based on heuristic transformations of existing data [26]. However, many augmentations used for natural images do not translate to the medical domain as the underlying data acquisition processes are different. This motivates the need for transformations that emulate the physics of medical image acquisition. Such specialized augmentations have been developed for computed tomography [19], magnetic resonance tomography [24] and ultrasound imaging [31]. Even though the physics behind OCTA image acquisition is well understood [27], data augmentation schemes specific to this modality have not been introduced to date.

3 Methods

3.1 Physiology-based simulation of the retinal vasculature

In order to generate diverse and realistic models of the retinal vasculature, we adapt and extend a physiology-based simulation by Schneider *et al.* [23]. In the following, we briefly describe their method before presenting our key novel contributions that allow us to model the highly unique vasculature of the retina (see figure 2).

In Schneider *et al.*'s simulation, blood vessels are represented as a forest of rooted tree graphs. Each graph node encodes a vessel segment modeled as a cylindrical tube with a specific radius, length and position as well as connectivity information to its parent vessel and up to two children. These trees are grown iteratively in a multi-scale fashion. At each simulation step, the concentration maps of oxygen and vascular endothelial growth factor (VEGF) are calculated depending on the spatial configuration of the vessel trees. High VEGF concentration, which drives the growth of new blood vessels, will primarily occur in

regions of low perfusion. To counteract the lack of oxygen, new vessels develop from nearby graph nodes via sprouting, elongation or bifurcation. Vessel formation adheres to rules derived from fluid dynamics that enforce morphologically plausible branching statistics.




Fig. 2. Main steps during the modeling of retinal vessel graphs. Additional examples can be found in the supplementary material.

Sequential simulation of the retinal vascular complexes The inner retina is supplied with blood by four vascular plexuses, which can be grouped into two complexes: the superficial vascular complex (SVC) and the deep vascular complex (DVC) [2,9]. We sequentially simulate these vascular complexes according to their order in blood stream direction.

The SVC features several large arterioles that branch out as they run towards the center of the retina. Its simulation is initialized with large vessel stumps that are radially distributed along the edges of the volume. The SVC leads to the DVC, which primarily consists of smaller, sprawling vessels. After the SVC’s simulation has converged, we calculate the chance of each vessel node to form a vertical sprout following the global branching statistics [23]. These sprouts are used as initial seeding vessels for the simulation of the DVC. We have heuristically tuned the simulations’ hyperparameters so that the generated vasculature models mimic the morphological statistics of the SVC and DVC, respectively. We provide the two sets of hyperparameters and rationale for these settings in the supplementary material.

Modeling of regions with low angiogenesis The outer retinal layers, which are responsible for the sensing of light, are not directly perfused. Furthermore, all retinal plexuses feature a region devoid of blood vessels at their center, the foveal avascular zone, so that light can pass through to the photoreceptor layers [30]. The lack of angiogenesis in these regions is induced by enforcing a reduced secretion of VEGF using binary masks. The shape of these masks is based on previously reported sizes for the retinal layers and the foveal avascular zone [22].

Deformation to the retinal layer structure We model both vascular complexes inside thin rectangular volumes before deforming the vessel graphs to the typical curved shape of the retina as it appears in OCT and OCTA images. To this end, we use the IOWA reference algorithms to extract retinal layer segmentations from anatomical OCT scans contained in the OCTA-500 dataset [1,5,10,12]. The SVC is sheared to adhere to the shape of the ganglion cell layer.

The DVC is deformed to the relief of the inner nuclear layer. Afterwards, we remodel the entire vascular tree one final time in order to enforce physiologically correct branching angles and vessel thinning along blood flow direction.

Conversion to gray-scale images By running the simulation with varying starting conditions, retinal layer shapes and random number generator seeds, we are able to obtain a diverse set of realistic retinal vessel graphs. The graphs are converted to voxelized gray-scale images with matching ground truth labels. For the following experiments we generate both two-dimensional *en-face* images (i.e. maximum intensity projections onto the axial plane) and three-dimensional image volumes.

3.2 Image augmentations based on the OCTA acquisition process

Next, we apply a series of newly developed image transformations that simulate image features and artifacts introduced during OCTA acquisition (see figure 3). We briefly describe these physics-based augmentations below, while providing pseudo-code and hyperparameters in the supplementary material.




Fig. 3. *En-face* OCTA images and close-up views of common artifacts in real OCT images from the ROSE dataset [17] and representative samples of our synthetically generated dataset. Additional examples can be found in the supplementary material.

Flow projection artifacts OCTA resolves the temporal variation of successive scans of the same volume. Blood cells flowing through the vessels will display as bright signal, while stationary tissue will appear dark. However, as the imaging beam traverses the retina, blood vessels will cast a shadow onto the deeper parts of the retina. This shadow is detected as temporally changing signal and wrongly appears as blood flow in the calculated angiograph [27]. We simulate such artifacts by identifying all vessels inside the volume that exceed a certain radius. Subsequently, we increase the image brightness along the projection of these vessels in the imaging beam’s direction.

Eye motion artifacts Volumetric OCTA scans are usually acquired in a rasterized fashion, in which multiple B-scans and small imaging patches are merged together in a postprocessing step. Movement of the subject’s eye or head during imaging may result in discontinuities between adjacent B-scans or so-called

quilting artifacts at the border of the patches [27]. We simulate four different motion artifacts along random horizontal and vertical cuts of the synthetic images: shearing, stretching, buckling and whiteout. The first three are modeled by shifting the intensity values on one side of the cut perpendicular to the imaging direction. The latter one is modeled as uniform noise along the cut and emulates a complete signal decorrelation in a single B-scan.

Vitreous floater artifacts During acquisition of retinal OCTA images, the light beam and its reflection pass through several transparent components of the eye, including the cornea, pupil, lens and vitreous. Opaque objects inside these structures, in particular vitreous floaters, can cause a loss of spatial coherence of the imaging beam that manifests itself as darkened areas in the image [27]. We model the typical thread-like shape of vitreous floaters as a graph of multiple connected line segments with random thickness, length and orientation. We determine the projection of the resulting polygon in imaging direction and reduce the signal in this region.

Background signal caused by capillary vessels At its smallest scale, the circulatory system branches into numerous capillary vessels. OCTA scanners are typically unable to resolve these small blood vessels, which appear as milky background signal in the image [28]. It is computationally infeasible to simulate the entire capillary vasculature, so that the background of our synthetic images appears unrealistically empty. To account for this limitation, we add binomial noise, which is blurred by a Gaussian filter.

4 Experiments and results

We evaluate the utility of our synthetic data in several benchmarking experiments using the publicly available ROSE-1 dataset [17]. We train supervised vessel segmentation algorithms using either synthetic OCTA image-label pairs ("synthetic") or real OCTA images with expert-derived labels ("real"). For all experiments we assess the algorithms' performance on real OCTA images by calculating accuracy (ACC), area under the receiver operating characteristic curve (AUC), F1-score (DICE) and the topology-aware centerline DICE (cLDICE) [25]. We adapt network implementations, data preprocessing and hyperparameter settings from Ma *et al.*⁵. Diverging from their approach that uses a fixed number of training epochs, we introduce an independent validation split for model selection. As a result the absolute performance and relative rankings of the vary algorithms may slightly differ compared to their reported results. All results are reported as mean and standard deviation across five cross-validation splits.


We find that the use of synthetic data facilitates vessel segmentation in OCTA images across all four considered neural network architectures. However, the quantitative scores are consistently lower than those of networks that were trained on the same domain (see table 1). Qualitatively, we find our segmentations to be superior (see figure 4). We suspect that this discrepancy is caused

⁵ <https://github.com/iMED-Lab/OCTA-Net-OCTA-Vessel-Segmentation-Network>

Table 1. Performance of four segmentation networks trained on either synthetic image-label pairs or real images with expert-derived annotations.

Model	Training data	ACC	AUC	DICE	clDICE
U-Net [21]	Real	0.89±0.01	0.91±0.01	0.69±0.02	0.60±0.02
	Synthetic	0.80±0.02	0.81±0.01	0.54±0.02	0.47±0.02
ResU-Net [34]	Real	0.88±0.01	0.90±0.01	0.68±0.02	0.59±0.02
	Synthetic	0.84±0.01	0.82±0.02	0.56±0.03	0.48±0.04
CS-Net [18]	Real	0.89±0.01	0.90±0.01	0.69±0.01	0.59±0.02
	Synthetic	0.77±0.04	0.81±0.02	0.55±0.02	0.48±0.02
CE-Net [8]	Real	0.88±0.01	0.88±0.01	0.67±0.02	0.58±0.02
	Synthetic	0.85±0.02	0.84±0.02	0.59±0.02	0.51±0.04

by a systematic difference in the datasets’ annotations. The expert labels in the ROSE-1 dataset do not encompass the smallest capillaries, whereas these vessels are labeled in our synthetic dataset. Consequently, the networks trained on synthetic data label these important small vessels, whereas the baselines do not.

**Fig. 4.** Two representative test samples with segmentation maps overlaid in orange. Algorithms trained on synthetic data accurately segment capillary vessels.

An advantage of our synthetic data pipeline is that it can easily generate large amounts of training data. In additional experiments, we show that network performance increases with training dataset size (see table 2). We also observe that top-end network performance improves when pretraining networks on synthetic data before finetuning them on a small number of real samples. In further ablation experiments, we quantified the benefit of both newly introduced components. Using vessel trees without the typical structure of the SVC and DVC results in a decreased segmentation accuracy (see table 3). Similarly, we find that leaving out the physics-based image transformations and training only with generic augmentations reduces performance.


Table 2. U-Net performance across different dataset types and training strategies.

Training data	Dataset size	ACC	AUC	DICE	clDICE
Synthetic	32	0.80±0.02	0.81±0.01	0.54±0.02	0.47±0.02
Synthetic	320	0.81±0.02	0.83±0.01	0.56±0.01	0.48±0.02
Synthetic	3200	0.81±0.01	0.84±0.01	0.57±0.01	0.49±0.02
Synthetic + finetuning	352	0.91±0.01	0.91±0.01	0.71±0.01	0.62±0.02
Real	32	0.89±0.01	0.91±0.01	0.69±0.02	0.60±0.02

Table 3. U-Net performance after ablating the physiology-based simulation of the retinal vasculature or the physics-based augmentations of our data synthesis method.

Configuration	ACC	AUC	DICE	clDICE
Proposed method	0.81±0.02	0.83±0.01	0.56±0.01	0.48±0.02
–vessel simulation	0.79±0.01	0.80±0.01	0.53±0.01	0.45±0.01
–augmentations	0.79±0.01	0.70±0.02	0.52±0.01	0.47±0.01

Finally, we train a 3D U-Net for vessel segmentation in volumes of the OCTA-500 dataset using only synthetic training data [11,12]. As of now, we are unable to benchmark on this task as there are OCTA datasets with volumetric labels are not available. Instead, we present initial qualitative results (see figure 5). These highlight the potential of our method to facilitate three-dimensional analysis of OCTA images; a task with enormous diagnostic potential which can only be rarely conducted due to a lack of suitable tools [33].

**Fig. 5.** Three-dimensional vessel segmentation of a volumetric OCTA image by a U-Net trained exclusively on our synthetic data.

5 Discussion and conclusion

In this work we have presented a method to generate highly realistic, synthetic OCTA images. Our method models the retinal vasculature with excellent anatomical detail using a physiology-based angiogenesis model. Moreover, we simulate image features and artifacts resulting from the OCTA acquisition process using a set of physics-based image augmentations. While our synthetic

OCTA images are still distinguishable from real images, they are accompanied by a fully characterized vascular graph. This intrinsically matched ground truth is highly advantageous when using the dataset for downstream tasks.

We demonstrate the promise of our approach by successfully training several segmentation algorithms without manual ground-truth annotations. In the future, our method can be tailored to emulate the characteristics of diverse clinical datasets. The vasculature simulation can be adapted to model the retinal morphology of specific populations, such as older subjects or patients with ocular, cardiac or neurological pathologies [9]. The image augmentations can be tuned to the hard- and software build into different OCTA scanners [28]. We believe that this versatility will enable application of our method beyond vessel segmentation and ultimately advance the quantitative analysis of OCTA in clinical practice.

Acknowledgments

Johannes C. Paetzold is supported by the DCoMEX project, financed by the Federal Ministry of Education and Research of Germany. Benjamin Knier is funded by the Else Kröner-Fresenius-Stiftung, the Gemeinnützige Hertie Foundation and received a research award from Novartis.

References

1. Abramoff, M.D., Garvin, M.K., Sonka, M.: Retinal imaging and image analysis. *IEEE reviews in biomedical engineering* **3**, 169–208 (2010)
2. Campbell, J., et al.: Detailed vascular anatomy of the human retina by projection-resolved optical coherence tomography angiography. *Scientific reports* **7**(1), 1–11 (2017)
3. Frid-Adar, M., Amer, R., Greenspan, H.: Endotracheal tube detection and segmentation in chest radiographs using synthetic data. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 784–792. Springer (2019)
4. Garcia-Garcia, A., Orts-Escolano, S., Oprea, S., Villena-Martinez, V., Martinez-Gonzalez, P., Garcia-Rodriguez, J.: A survey on deep learning techniques for image and video semantic segmentation. *Applied Soft Computing* **70**, 41–65 (2018)
5. Garvin, M.K., Abramoff, M.D., Wu, X., Russell, S.R., Burns, T.L., Sonka, M.: Automated 3-D intraretinal layer segmentation of macular spectral-domain optical coherence tomography images. *IEEE transactions on medical imaging* **28**(9), 1436–1447 (2009)
6. Gerl, S., et al.: A distance-based loss for smooth and continuous skin layer segmentation in optoacoustic images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 309–319. Springer (2020)
7. Gherardini, M., Mazomenos, E., Mencias, A., Stoyanov, D.: Catheter segmentation in X-ray fluoroscopy using synthetic data and transfer learning with light U-nets. *Computer Methods and Programs in Biomedicine* **192**, 105420 (2020)
8. Gu, Z., et al.: CE-Net: Context encoder network for 2D medical image segmentation. *IEEE transactions on medical imaging* **38**(10), 2281–2292 (2019)

9. Joussen, A.M., Gardner, T.W., Kirchhof, B., Ryan, S.J.: Retinal vascular disease, vol. 75. Springer (2007)
10. Li, K., Wu, X., Chen, D.Z., Sonka, M.: Optimal surface segmentation in volumetric images—a graph-theoretic approach. *IEEE transactions on pattern analysis and machine intelligence* **28**(1), 119–134 (2005)
11. Li, M., et al.: Image projection network: 3D to 2D image segmentation in OCTA images. *IEEE Transactions on Medical Imaging* **39**(11), 3343–3354 (2020)
12. Li, M.o.: IPN-V2 and OCTA-500: Methodology and dataset for retinal image segmentation. *arXiv preprint arXiv:2012.07261* (2020)
13. Liew, G., Wang, J.J.: Retinal vascular signs: a window to the heart? *Revista Española de Cardiología (English Edition)* **64**(6), 515–521 (2011)
14. Lindner, L., Narnhofer, D., Weber, M., Gsaxner, C., Kolodziej, M., Egger, J.: Using synthetic training data for deep learning-based GBM segmentation. In: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). pp. 6724–6729. IEEE (2019)
15. Liu, Y., et al.: Variational intensity cross channel encoder for unsupervised vessel segmentation on OCT angiography. In: *Medical Imaging 2020: Image Processing*. vol. 11313, p. 113130Y. International Society for Optics and Photonics (2020)
16. London, A., Benhar, I., Schwartz, M.: The retina as a window to the brain—from eye research to CNS disorders. *Nature Reviews Neurology* **9**(1), 44–53 (2013)
17. Ma, Y., et al.: ROSE: a retinal OCT-angiography vessel segmentation dataset and new model. *IEEE transactions on medical imaging* **40**(3), 928–939 (2020)
18. Mou, L., et al.: CS-Net: channel and spatial attention network for curvilinear structure segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 721–730. Springer (2019)
19. Omigbodun, A.O., Noo, F., McNitt-Gray, M., Hsu, W., Hsieh, S.S.: The effects of physics-based data augmentation on the generalizability of deep neural networks: Demonstration on nodule false-positive reduction. *Medical physics* **46**(10), 4563–4574 (2019)
20. Paetzold, J.C., et al.: Transfer learning from synthetic data reduces need for labels to segment brain vasculature and neural pathways in 3d. In: *International Conference on Medical Imaging with Deep Learning—Extended Abstract Track* (2019)
21. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
22. Samara, W.A., et al.: Correlation of foveal avascular zone size with foveal morphology in normal eyes using optical coherence tomography angiography. *Retina* **35**(11), 2188–2195 (2015)
23. Schneider, M., Reichold, J., Weber, B., Székely, G., Hirsch, S.: Tissue metabolism driven arterial tree generation. *Medical image analysis* **16**(7), 1397–1414 (2012)
24. Shaw, R., Sudre, C.H., Varsavsky, T., Ourselin, S., Cardoso, M.J.: A k-space model of movement artefacts: application to segmentation augmentation and artefact removal. *IEEE transactions on medical imaging* **39**(9), 2881–2892 (2020)
25. Shit, S., et al.: cDice - a novel topology-preserving loss function for tubular structure segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16560–16569 (2021)
26. Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. *Journal of big data* **6**(1), 1–48 (2019)
27. Spaide, R.F., Fujimoto, J.G., Waheed, N.K.: Image artifacts in optical coherence angiography. *Retina (Philadelphia, Pa.)* **35**(11), 2163 (2015)

28. Spaide, R.F., Fujimoto, J.G., Waheed, N.K., Sadda, S.R., Staurengi, G.: Optical coherence tomography angiography. *Progress in retinal and eye research* **64**, 1–55 (2018)
29. Spaide, R.F., Klancnik, J.M., Cooney, M.J.: Retinal vascular layers imaged by fluorescein angiography and optical coherence tomography angiography. *JAMA ophthalmology* **133**(1), 45–50 (2015)
30. Tick, S., et al.: Foveal shape and structure in a normal population. *Investigative ophthalmology & visual science* **52**(8), 5105–5110 (2011)
31. Tirindelli, M., Eilers, C., Simson, W., Paschali, M., Azampour, M.F., Navab, N.: Rethinking ultrasound augmentation: A physics-inspired approach. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 690–700. Springer (2021)
32. Todorov, M.I., et al.: Machine learning analysis of whole mouse brain vasculature. *Nature methods* **17**(4), 442–449 (2020)
33. Yu, S., et al.: Cross-domain depth estimation network for 3D vessel reconstruction in OCT angiography. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 13–23. Springer (2021)
34. Zhang, Z., Liu, Q., Wang, Y.: Road extraction by deep residual U-net. *IEEE Geoscience and Remote Sensing Letters* **15**(5), 749–753 (2018)

Supplementary material




Fig. 6. Random examples of synthetically generated retinal vessel graphs.

Table 4. Tuned hyperparameters for the simulation of retinal vessel graphs. The changes compared to the implementation by Schneider *et al.* (original) facilitate the synthesis of the superior vascular complex (SVC) and deep vascular complex (DVC).

Parameter	original	SVC	DVC	Rationale behind change
Ω_x	936 μm	3200 μm	3200 μm	Flat simulation volume for modeling of retinal vasculature
Ω_y	864 μm	3200 μm	3200 μm	Flat simulation volume for modeling of retinal vasculature
Ω_z	1600 μm	64 μm	64 μm	Flat simulation volume for modeling of retinal vasculature
θ_c	1.050	1.025	1.025	Encourage more sprouting of vessels
θ_p	0.95	0.90	0.90	Faster model convergence without noticeable loss of detail
r_{initial}	15.0 μm	22.5 μm	–	Larger initial radius of the retinal arterioles
r_{min}	3.0 μm	2.25 μm	1.5 μm	Simulation of smallest retinal capillaries
r_{degen}	5.00 μm	3.75 μm	2.50 μm	Simulation of smallest retinal capillaries
r_{prune}	2.00 μm	1.5 μm	1.00 μm	Simulation of smallest retinal capillaries
m_b	14	16	12	Longer/ shorter vessel segments before bifurcation
λ_g	1.00	1.00	0.25	Contorted, web-like structure of DVC
γ	3.0	3.0	2.5	Larger bifurcation angles in DVC




Fig. 7. Three additional examples of artifacts in real and synthetic OCTA images.

Pseudo-code for the physics-based image augmentations. The default hyperparameters used in this study are included in the algorithms' requirements. Units of size are provided in pixels/ voxels unless explicitly specified.

Algorithm 1 Vitreous floater artifacts

Require: Image I , empty grid of image shape in axial plane F , segment number distribution $S : \mathcal{U}(10, 20)$, length distribution $\mathcal{L} : \mathcal{N}(30, 10)$, thickness distribution $\mathcal{T} : \mathcal{N}(10, 5)$, floater opacity distribution $\mathcal{Q} : \mathcal{U}(0.5, 1.0)$, Gaussian kernel standard deviation $\nu : 10$

- 1: $p \leftarrow \text{SAMPLE STARTING POSITION}(F)$ ▷ Random pixel on grid
- 2: $s \leftarrow \text{SAMPLE}(S)$ ▷ Number of segments
- 3: **for all** s segments **do**
- 4: $d, l, t \leftarrow \text{SAMPLE SEGMENT GEOMETRY}(\mathcal{L}, \mathcal{T})$ ▷ Direction, length and thickness of segment
- 5: $F \leftarrow \text{DRAW LINE SEGMENT}(F, p, d, l, t)$ ▷ Add to floater canvas
- 6: $p \leftarrow p + l \cdot d$
- 7: $F \leftarrow \text{GAUSSIAN BLURRING}(F, \nu)$
- 8: $P \leftarrow \text{CREATE PROJECTION}(F)$ ▷ Shadow in imaging direction
- 9: $q \leftarrow \text{SAMPLE}(\mathcal{Q})$ ▷ Floater opacity
- 10: $I \leftarrow I * q \cdot F$
- 11: **return** I

Algorithm 2 Flow projection artifacts

Require: Image I , Forest of vessel graphs \mathcal{G} , minimum radius $r : 10 \mu\text{m}$, projection scale $s : 0.5$

- 1: $\mathcal{G}' \leftarrow \text{REMOVE SMALL VESSELS}(\mathcal{G}, r)$ ▷ Only vessels of a certain size are considered
- 2: $J \leftarrow \text{VOXELIZE}(\mathcal{G}')$
- 3: $P \leftarrow \text{CREATE PROJECTION}(J)$ ▷ Binary mask of shadow in imaging direction
- 4: $I \leftarrow I + s \cdot P$
- 5: $I \leftarrow \text{CLIP}(I)$ ▷ Clip to maximum and minimum intensity
- 6: **return** I

Algorithm 3 Motion artifacts in vertical direction (equivalent in horizontal direction)

Require: Image I , cut number distribution $C : \mathcal{U}(2, 5)$, shift distribution $S : \mathcal{U}(0, 5)$

- 1: $c \leftarrow \text{SAMPLE}(C)$ ▷ Number of cuts
- 2: **for all** c cuts **do**
- 3: $t \leftarrow \text{SAMPLE MOTION TYPE}$ ▷ Up, down, left, right or whiteout
- 4: **if** t is not whiteout **then**
- 5: $s \leftarrow \text{SAMPLE}(S)$ ▷ Magnitude of motion
- 6: $I_l, I_r \leftarrow \text{Split image}(I)$ ▷ Cut image vertically and return both parts
- 7: $I'_r \leftarrow \text{GET SHIFTED IMAGE PART}(I, I_r, t, s)$ ▷ Get shifted region from original image
- 8: $I \leftarrow \text{MERGE}(I_l, I'_r)$
- 9: **else** ▷ Whiteout
- 10: $I_c \leftarrow \text{CUT IMAGE}$ ▷ Return column along cut
- 11: $N \leftarrow \text{SAMPLE}(\mathcal{U}(0.5, 1.0))$ ▷ Uniform noise in shape of cut
- 12: $I \leftarrow \text{REPLACE}(I, I_c, N)$ ▷ Replace cut with noise
- 13: **return** I

Algorithm 4 Background signal caused by capillary vessels

Require: Image I , noise rate $p : 0.01$, noise dilation steps $d : 1$, noise scale $s : 0.33$

- 1: $N \leftarrow \text{SAMPLE}(\mathcal{B}(1, p))$ ▷ Binomial noise of image shape
- 2: $N \leftarrow \text{BINARY DILATION}(N, d)$
- 3: $I \leftarrow (1 - s) \cdot I + s \cdot N$
- 4: **return** I
