

A Label-Free and Data-Free Training Strategy for Vasculature Segmentation in serial sectioning OCT Data

Etienne Chollet
Yaël Balbastre
Caroline Magnain
Bruce Fischl*
Hui Wang*

ECHOLLET@MGH.HARVARD.EDU
YBALBASTRE@MGH.HARVARD.EDU
CMAGNAIN@MGH.HARVARD.EDU
BFISCHL@MGH.HARVARD.EDU
HWANG47@MGH.HARVARD.EDU

A.A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Boston, USA

Abstract

Serial sectioning Optical Coherence Tomography (sOCT) is a high-throughput, label free microscopic imaging technique that is becoming increasingly popular to study post-mortem neurovasculature. Quantitative analysis of the vasculature requires highly accurate segmentation; however, sOCT has low signal-to-noise-ratio and displays a wide range of contrasts and artifacts that depend on acquisition parameters. Furthermore, labeled data is scarce and extremely time consuming to generate. Here, we leverage synthetic datasets of vessels to train a deep learning segmentation model. We construct the vessels with semi-realistic splines that simulate the vascular geometry and compare our model with realistic vascular labels generated by constrained constructive optimization. Both approaches yield similar Dice scores, although with very different false positive and false negative rates. This method addresses the complexity inherent in OCT images and paves the way for more accurate and efficient analysis of neurovascular structures.

Keywords: Medical Imaging, Segmentation, Optical coherence tomography, Machine Learning, Data Synthesis.

1. Introduction

Vascular networks are central to brain function in health and disease, but mapping the human neurovasculature across scales, from mesoscale veins and arteries down to microscale capillaries, is an unmet challenge. Optical Coherence Tomography (OCT) (Huang et al., 1991) offers a wealth of information concerning vascular details. Combined with serial sectioning, sOCT reconstructs volumetric microvascular networks (Wang et al., 2018), but their automated segmentation is hindered by the presence of complex, structured, high-frequency noise. Traditional knowledge-based methods, that rely on hessian-based filters (Frangi et al., 1998; Yang et al., 2022), morphological operations (Zana and Klein, 2001), or region growing techniques (Dokl adal et al., 1999) are highly sensitive to noise and fall short especially when vessels cross tissue boundaries, or when examining non-tubular morphologies. Conversely, convolutional neural networks (CNNs) have proven robust to these inconsistencies (Goni et al., 2022). However, the cost of producing labeled data results in low generalizability of the CNN models due to a lack of diverse training data. As an alternative, synthetic datasets have been increasingly leveraged in tasks of segmentation and

* Joint senior authors

registration (Hoffmann et al., 2021; Billot et al., 2023). Synthetic vascular labels generated by constrained constructive optimization (CCO) (Schneider et al., 2012) have been used to train segmentation models for retinal OCT (Kreitner et al., 2024) and photoacoustic images (Sweeney et al., 2023), and to pretrain models targeting MR angiography and light-sheet microscopy (Paetzold et al., 2019; Tetteh et al., 2020). In contrast to realistic synthesis, Dey et al. (2024) recently showed that simple geometrical analogy can effectively yield a universal segmentation model in cells. Following this line of thought, we:

1. Propose a domain-randomized image generation strategy for sOCT volumes;
2. Propose an *unrealistic* spline-based synthesis engine that covers a wide range of vessel-like geometries, much larger than seen in practice;
3. Compare models trained with realistic CCO labels and unrealistic spline labels.




Figure 1: 3D renderings of CCO (A) and complex (B) labels, and (C) data synthesis pipeline for unsupervised training of UNet in vasculature segmentation task.

2. Methods

Spline-based label synthesis. We generate vascular trees by randomly drawing cubic splines within a 3D space. We initially sample straight lines before adding randomly jittered control points. Their number and the magnitude of the jitter define the spline tortuosity. Branching points are then sampled to serve as the endpoints of a second set of splines, and the procedure is repeated recursively. The spline radius is allowed to vary along its length and is also encoded by cubic splines. The change in radius across levels in the recursion is randomly sampled. Finally, splines are rasterized using a distance transform computed by Gauss-Newton optimization. We generated two different datasets, labeled *complex* and *simple*, that respectively used wide and peaked distribution for all hyper-parameters.

CCO-based label synthesis. We used the synthetic labels released in Tetteh et al. (2020), which were generated using the procedure described in Schneider et al. (2012).

Domain-randomized OCT synthesis. Vessel-specific and parenchyma-specific textures are generated by sampling two different random label maps, following the procedure from Hoffmann et al. (2021). A random intensity, between 0 and 1, is assigned to each label, and the parenchyma and vessels are fused *via* voxel-wise multiplication. Thus, a vessel’s intensity is always lower than its surrounding tissue. (*i.e.*, "dark vessels"). This is crucial in OCT, as refractive effects can generate an artifactual bright shadow above the true vessel. Finally, random slab-wise intensity non-uniformities (that cover absorption, depth of focus,

stitching artifacts between slices) and speckle noise are applied.

Model and training. A U-Net (Ronneberger et al., 2015) with residual blocks (He et al., 2016) (number of parameters: $7.4 \cdot 10^6$) and Tilborghs et al. (2022)’s Dice loss was trained with Adam (learning rate: 10^{-2} with 2,000 warmup steps and 20,000 cooldown steps) for 100,000 steps. Each model was trained four times with different initial random weights. All training patches were of size $128 \times 128 \times 128$.

Validation. All models were evaluated on a $301 \times 301 \times 301$ voxel OCT volume manually labeled by an expert. A sine-weighted moving window was used at inference time.

A schematic of the complete pipeline is presented in figure 1.




Figure 2: (A): Qualitative comparison of model performance. (B): Dice score, false positive rate (FPR), and false negative rate (FNR) for models (n=4) trained on complex, CCO, and simple datasets. ANOVA analysis was used to delineate statistical significance (****: $P < 0.0001$, ***: $P < 0.001$)

3. Results & Discussion

Figure 2 presents a comparison of the mean dice scores, false positive rates (FPR), and false negative rates (FNR) for each synthesis pathway. A two-way ANOVA shows that while vessel labels with simple parameters hurt network performance, networks trained on complex labels (dice=0.627) and CCO labels (dice=0.605) were similarly accurate. However, these two networks have drastically different error patterns, with the CCO models tending to undersegment (FPR=0.055, FNR=0.511) and the complex models tending to oversegment (FPR=0.330, FNR=0.213). A qualitative assessment of the segmentation shows that many of the false positives are in fact small but real vessels that were missed by the expert (see *e.g.* figure 2(A.b)). This is the foundation of the benefit of label-conditioned image generation, which guarantees a precise alignment of labels and image attributes. To train a segmentation model, we have shown that it is unnecessary to construct vessels according to rigid physical models. This discovery highlights the possibility of using synthetic data to build accurate segmentation models.

Acknowledgments

Much of the computation resources required for this research was performed on computational hardware generously provided by the [Massachusetts Life Sciences Center](#).

References

- Benjamin Billot, Douglas N Greve, Oula Puonti, Axel Thielscher, Koen Van Leemput, Bruce Fischl, Adrian V Dalca, Juan Eugenio Iglesias, et al. SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Medical image analysis*, 86:102789, 2023.
- Neel Dey, Mazdak Abulnaga, Benjamin Billot, Esra Abaci Turk, Ellen Grant, Adrian V Dalca, and Polina Golland. Anystar: Domain randomized universal star-convex 3D instance segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7593–7603, 2024.
- Petr Dokládal, Christophe Lohou, Laurent Perroton, and Gilles Bertrand. Liver blood vessels extraction by a 3-D topological approach. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI’99: Second International Conference, Cambridge, UK, September 19-22, 1999. Proceedings 2*, pages 98–105. Springer, 1999.
- Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. Multiscale vessel enhancement filtering. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI’98: First International Conference Cambridge, MA, USA, October 11–13, 1998 Proceedings 1*, pages 130–137. Springer, 1998.
- Mohammad Raihan Goni, Nur Intan Raihana Ruhaiyem, Muzaimi Mustapha, Anusha Achuthan, and Che Mohd Nasril Che Mohd Nassir. Brain vessel segmentation using deep learning—a review. *IEEE Access*, 10:111322–111336, 2022.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 630–645. Springer, 2016.
- Malte Hoffmann, Benjamin Billot, Douglas N Greve, Juan Eugenio Iglesias, Bruce Fischl, and Adrian V Dalca. SynthMorph: learning contrast-invariant registration without acquired images. *IEEE transactions on medical imaging*, 41(3):543–558, 2021.
- David Huang, Eric A Swanson, Charles P Lin, Joel S Schuman, William G Stinson, Warren Chang, Michael R Hee, Thomas Flotte, Kenton Gregory, Carmen A Puliafito, et al. Optical coherence tomography. *science*, 254(5035):1178–1181, 1991.
- Linus Kreitner, Johannes C Paetzold, Nikolaus Rauch, Chen Chen, Ahmed M Hagag, Alaa E Fayed, Sobha Sivaprasad, Sebastian Rausch, Julian Weichsel, Bjoern H Menze, et al. Synthetic optical coherence tomography angiographs for detailed retinal vessel segmentation without human annotations. *IEEE Transactions on Medical Imaging*, 2024.

- Johannes C Paetzold, Oliver Schoppe, Rami Al-Maskari, Giles Tetteh, Velizar Efremov, Mihail I Todorov, Ruiyao Cai, Hongcheng Mai, Zhouyi Rong, Ali Ertuerk, et al. Transfer learning from synthetic data reduces need for labels to segment brain vasculature and neural pathways in 3d. 2019.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- Matthias Schneider, Johannes Reichold, Bruno Weber, Gábor Székely, and Sven Hirsch. Tissue metabolism driven arterial tree generation. *Medical image analysis*, 16(7):1397–1414, 2012.
- Paul W Sweeney, Lina Hacker, Thierry L Lefebvre, Emma L Brown, Janek Gröhl, and Sarah E Bohndiek. Segmentation of 3D blood vessel networks using unsupervised deep learning. *bioRxiv*, pages 2023–04, 2023.
- Giles Tetteh, Velizar Efremov, Nils D Forkert, Matthias Schneider, Jan Kirschke, Bruno Weber, Claus Zimmer, Marie Piraud, and Björn H Menze. Deepvesselnet: Vessel segmentation, centerline prediction, and bifurcation detection in 3-d angiographic volumes. *Frontiers in Neuroscience*, 14:592352, 2020.
- Sofie Tilborghs, Jeroen Bertels, David Robben, Dirk Vandermeulen, and Frederik Maes. The dice loss in the context of missing or empty labels: introducing ϕ and ϵ . In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 527–537. Springer, 2022.
- Hui Wang, Caroline Magnain, Ruopeng Wang, Jay Dubb, Ani Varjabedian, Lee S Tirrell, Allison Stevens, Jean C Augustinack, Ender Konukoglu, Iman Aganj, et al. as-PSOCT: Volumetric microscopic imaging of human brain architecture and connectivity. *Neuroimage*, 165:56–68, 2018.
- Jiarui Yang, Shuaibin Chang, Ichun Anderson Chen, Sreekanth Kura, Grace A Rosen, Nicole A Saltiel, Bertrand R Huber, Divya Varadarajan, Yael Balbastre, Caroline Magnain, et al. Volumetric characterization of microvasculature in ex vivo human brain samples by serial sectioning optical coherence tomography. *IEEE Transactions on Biomedical Engineering*, 69(12):3645–3656, 2022.
- Frederic Zana and J-C Klein. Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE transactions on image processing*, 10(7):1010–1019, 2001.