Robust automatic brain vessel segmentation in 3D CTA scans using dynamic 4D-CTA data

Alberto Mario Ceballos Arroyo[†,*,1], Shrikanth M. Yadav[†,2,3,5], Chu-Hsuan Lin[4,5], Jisoo Kim[4,5], Geoffrey S. Young[4,5], Lei Qin[‡,3,5], and Huaizu Jiang [‡,1]

[1] Northeastern University, Boston, USA
[2] Washington University in St. Louis, St, Louis, USA
[3] Dana-Farber Cancer Institute, Boston, USA
[4] Brigham and Women's Hospital, Boston, USA
[5] Harvard Medical School, Boston, USA

Abstract. Existing datasets for brain vessel segmentation in CTA are limited in scale and largely restricted to major vessels, leaving peripheral vasculature underrepresented. We address this data-level bottleneck by introducing DynaVessel, a richly annotated dynamic 4D-CTA dataset, alongside a novel annotation pipeline that exploits multi-phase contrast enhancement and background subtraction to reduce the effort required for expert manual annotation of vessels. This not only enlarges our dataset by 5-6 fold, but also provides a novel data-driven method for training deep learning models to detect arteries and veins despite the wide variation in vessel contrast encountered in single-phase CTA scans. DynaVessel comprises 110 training volumes from 29 patients and 165 test volumes from 14 patients. By comparing with two similarly-sized datasets for brain vessel segmentation in CTA, we demonstrate that DynaVessel-trained nnUNet models produce superior segmentations compared to the same model trained on single phase data, with an average mDC of 0.846/0.957 for arteries/veins, shorter average directed Hausdorff distance (0.304 mm/0.08 mm for arteries/veins) and improved topology sensitivity (0.877/0.974 for arteries/veins) in the TopBrain dataset. Code, model weights, and instructions to access the dataset are available online.

Keywords: Vessel segmentation · Dynamic 4D-CTA · Deep learning

1 Introduction

Dynamic 4D-computed tomography angiography (CTA) [15] provides excellent temporal vascular information by acquiring images throughout the passage of intravenous contrast at multiple time points. Nonetheless, single-phase CTA remains the predominant clinical technique for neurovascular imaging because

† Ceballos Arroyo, A. and Yadav, S. contributed equally to this work
‡ Lei, Q. and Jiang, H. had an equal supervision role in this work
* corresponding author: ceballosarroyo.a@northeastern.edu

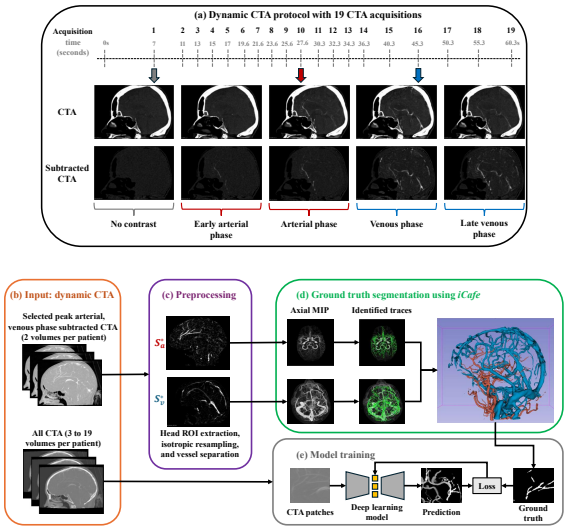

Fig. 1. Summary of our methodology: (a) Dynamic CTA are collected for all patients. (b) Peak arterial/venous phase scans are selected for semi-automatic annotation. (c) The subtracted images are processed to suppress voxels outside the head and resampled, then the artery-only and vein-only volumes are computed. (d) The vessel-separated images are processed using *iCafe* and manually validated. (e) For model training, all available CTA data (*i.e.*, 3-19 volumes per patient) are paired with the ground truths.

single phase scanners are far less expensive and far more widely available than multiphase capable scanners. Hence, robust methods to automatically extract the brain vasculature from single-phase CTA could have substantial clinical impact, both by directly assisting radiologists during interpretation of scans and by serving as an auxiliary input to image processing pipelines for tasks such as occlusion localization [10,13] and aneurysm detection [4,12].

Most work on CTA-based brain vessel segmentation has focused on the Circle of Willis (CoW). The TopCoW dataset [24] of 125 manually annotated CTA scans with CoW artery labels has had considerable impact [3,17], but provides no labels for clinically relevant vessels distal to the CoW. Two extensions address this: TopBrain [25], which adds veins and extra artery labels to a subset of TopCoW, and VesselVerse [1], which also expands TopCoW with annotations of veins and arteries (albeit unified as a binary vessel class). Some researchers have trained on large datasets to address this limitation; VesselFM [23] uses diverse modalities and extensive synthetic data. However, CTA comprises less than 1% of its training set and it has not been evaluated on head CTA, so the extent to which these approaches aid CTA intracranial vessel segmentation is unclear.

The limited anatomic coverage and/or scale of existing datasets reveals a data-level bottleneck for progress in brain vessel segmentation. Accurate annotation is challenging due to bone, soft tissue, and non-vascular contrast-enhanced structures requiring specialized anatomical expertise, which limits how quickly

new datasets can be produced. Rather than proposing a new segmentation architecture, this work directly addresses this bottleneck through a novel dataset and annotation pipeline. We exploit 4D-CTA's near-complete removal of bone and soft tissue and dynamic separation of arteries and veins to overcome this barrier (see Fig. 1), and we annotate scans from 43 patients with comprehensive artery and vein labels. Because of the absence of interphase motion, the labels obtained using our methodology can be re-used across all phases for the same patient. This allows us to train an nnUNet model [7] on multi-contrast data from 29 patients, achieving robust segmentation performance across contrast phases. In summary, we make the following contributions:

- We document a pipeline for annotating arteries and veins in head CTA scans based on dynamic 4D-CTA data, which significantly reduces annotation time
- We present DynaVessel, an intracranial vessel segmentation dataset comprising 275 head CTA scans from 43 patients with artery and vein annotations
- We evaluate a model trained on DynaVessel against several baselines on in-domain and out-of-domain benchmarks, achieving superior segmentation performance on both arteries and veins

2 Materials and methods

2.1 Data

This retrospective study received IRB approval with a waiver of informed consent (MGB HRC IRB Protocol #2022P000792). The 4D-CTA scans obtained from the Mass General Brigham (Boston, USA) clinical research database were acquired using a 320-detector-row CT scanner (Canon-Toshiba Aquilion One). Imaging was conducted in axial mode using the following parameters: 80 kV, 150 mA, 0.75 s rotation time, 320 slices with 0.5 mm thickness, and a 512 × 512 matrix. For each patient, image acquisition began roughly 7 seconds after injecting 75–100 mL of iodinated contrast agent (GE Omnipaque 350) at a rate of 4–5 mL/s, continuing intermittently through the baseline, arterial, and venous phases. As shown in Fig. 1(a), images were captured at 19 time points over a 60-second interval, where the peak arterial and venous phase images are of primary interest. The scanner uses the baseline image—acquired before contrast reaches the brain—to subtract bone and soft tissue from subsequent images. The 16 cm detector means patients experience little motion during the 60 second data acquisition, minimizing inter-phase patient motion. We also use intra-patient rigid registration (validated by a radiologist) between the peak arterial phase and other phases to ensure proper image alignment. Of 106 patients scanned in 2023, 29 had at least one arterial and one venous phase scan of acceptable image quality. Although these were acquired using dynamic 4D-CTA, complete phase datasets were not consistently archived in the PACS by the technicians. The available scans were annotated and used for model training. For evaluation, we retrieved data from 14 patients scanned between 2016 to 2022 who had dynamic CTA with at least one peak arterial and one peak venous phase image at acceptable resolution, with many also including more CTA phases.

Algorithm 1 Artery and vein suppression.

Input: S_a, S_v, X_a, X_v | **Output:** S_a^*, S_v^*
1: Compute rigid transformation G_{rv} such that X_a is moving and X_v is fixed
2: Compute rigid transformation G_{ra} such that X_v is moving and X_a is fixed
3: $S_{v\to a} \leftarrow G_{ra}(X_v)$, transform X_v so that it is in the same space as X_a
4: $S_{a\to v} \leftarrow G_{rv}(X_a)$, transform X_a so that it is in the same space as X_v
5: $s_a^* \leftarrow \begin{cases} 0 & \text{if } s_a < s_{v\to a} \\ s_a & \text{if } s_a > s_{v\to a} \end{cases}$, where s_a^*, $s_{v\to a}$, and s_a are elements of S_a^*, $S_{v\to a}$, and S_a
6: $s_v^* \leftarrow \begin{cases} 0 & \text{if } s_v < s_{a\to v} \\ s_v & \text{if } s_v > s_{a\to v} \end{cases}$, where s_v^*, $s_{a\to v}$, and s_v are elements of S_v^*, $S_{a\to v}$, and S_v

To suppress extracranial soft tissues, a binary ROI encompassing the skull and neck was delineated on a publicly available CT template [19]. The template was affine registered to the patient's CTA image, and the resulting transformation matrix was applied to the ROI to create a binary mask of the skull and neck in the patient's space. Voxels outside of this ROI were excluded from further analysis. All images were then resampled to an isotropic resolution of 0.468 mm. Let X_a be the arterial phase CTA, X_v the venous phase CTA, and S_a / S_v the soft-tissue-subtracted versions of X_a and X_v, Algorithm 1 was applied to the subtracted CTA images to separate arterial and venous structures into distinct volumes for each patient. The output of Algorithm 1 for each patient is two volumes: S_a^* (background- and vein-suppressed arterial phase) and S_v^* (background- and artery-suppressed venous phase).

Next, we label vessels with the *iCafe* [5] feature extraction pipeline. This interactive tool is designed for vessel annotation in MRA images, which have better arterial contrast than CTA. Nevertheless, we found that vessel-separated volumes S_a^* and S_v^* derived from 4D-CTA provide near-MRA levels of contrast, enabling us to use *iCafe* for both artery and vein delimitation. The tool uses intensity normalization followed by segmentation using the Phansalkar threshold [14] for vessel tracking and the Renyi entropy threshold [9] for visualization. A vessel tracing algorithm, based on a open-curve active contour model, is used to identify vessel-like structures [22]. Then, the trace and its corresponding mask are exported as a 3D segmentation of the vessels. The segmentations were imported into 3D Slicer [11] and co-registered with the subtracted S_a, S_v, and source CTA images. A fellowship-trained neuroradiologist with 5-years of subspecialty experience (JK) then performed slice-by-slice quality assessment and refinement of the segmentation masks, correcting false-positive and false-negative regions to generate the final vessel annotations. Fig. 1(d) shows the overall process of extracting the ground truth (GT), which results in a two-class (artery, vein) segmentation dataset, as depicted in Fig. 2.

2.2 Training and evaluation

Our training and evaluation data includes multiple CTA phases from the same patient, each paired with a corresponding GT segmentation. All images were rigidly registered to the arterial phase CTA using ANTS [20]. We split the train-

ing and test datasets at the patient level to prevent any patient from appearing in both sets. To train the nnUNet model, we use the nnUNet v2 library [7,8]. Aside from specifying the model type to be a Residual Encoder UNet (large preset), all models were trained with the combined dice + cross-entropy loss for 1000 epochs on a linearly decaying learning rate starting at 0.01, using the nnUNet planner's automatically selected hyper-parameters. All experiments were performed on a workstation with an NVIDIA RTX 4090 GPU, 8 CPU cores, and 64 GB of RAM.

DL models trained on our dataset output a volume with two classes (artery and vein), and we analyze how the predictions align with different vascular class to provide insight into region-specific performance. Specifically, let volume A_i represent the GT vessel mask of interest, where $i = 1, 2$ corresponds to arterial and venous labels; let volume P denote the predicted mask for either region. In line with established guidelines [24], we use the following metrics:

1. **Modified Dice coefficient**: we propose a modified version of the Dice coefficient (mDC) focusing on sensitivity, defined as:

$$mDC(A_i, P) = \frac{|A_i \cap P|}{|A_i|}$$

2. **Average directed Hausdorff surface distance**: we chose this variation of the Hausdorff distance based on previous work [2]. It is defined as:

$$adHD(A_{s,i}, P_s) = \frac{1}{|A_{s,i}|} \sum_{a \in A_{s,i}} \min_{p \in P_s} d(a, p)$$

where $A_{s,i}$ and P_s denote the surfaces of volumes A_i and P, respectively. $d(a, p)$ represents the Euclidean distance between surface points. $adHD$ severely penalizes cases where the boundaries of the predictions are far from the targets, as it sums minimum distances from each GT-surface point to the nearest predicted surface point.

3. **Topology sensitivity**: we define $tSens$, based on *clDice* [18] to quantify the extent of centerlines that are captured by the vessel segmentation models:

$$tSens(A_{c,i}, P) = \frac{|A_{c,i} \cap P|}{|A_{c,i}|}$$

where $A_{c,i}$ is the centerline volume of A_i.

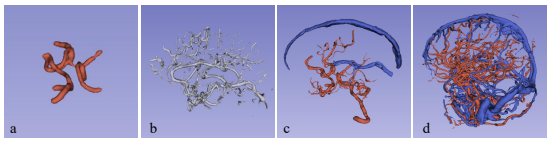

Fig. 2. Comparison of vascular structure coverage across datasets, with arteries in red and veins in blue: (a) TopCoW, (b) VesselVerse, (c) TopBrain, (d) our DynaVessel.

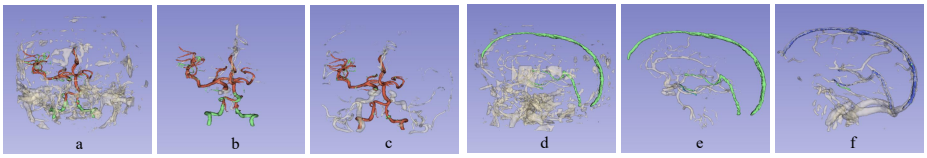

Fig. 3. Segmentations for a TopBrain scan. Left group (a, b, c): arteries. Right group (d, e, f): veins. (a, d): VesselFM; (b, e) VesselVerse nnUNet; (c, f): M_{DV}. True positive (TP) areas are **red** (arteries) or **blue** (veins), false positives (FP) are gray, and false negatives (FN) are **green**. Both baselines output a single vessel class, so artery predictions show up in the vein visualization as FPs, and vice-versa.

Table 1. Binary vessel segmentation and artery/vein segmentation results. Results that are statistically significantly worse compared to M_{DV} based on Wilcoxon signed-rank testing ($p < 0.05$) are highlighted with an asterisk.

Dataset	Area	Model	Train Set	mDC ↑	tSens ↑	adHD ↓
TopCoW ISLES	CoW	DynUNet	VesselFM	0.733±0.18*	0.865±0.15*	0.569±0.85*
		nnUNet	TopCoW	0.756±0.05*	0.921±0.05*	0.366±0.21*
		nnUNet	VesselV.	0.933±0.04	0.988±0.02	0.064±0.04
		nnUNet	DynaV.	**0.966**±0.03	**0.991**±0.02	**0.034**±0.02
VesselV.	Vessel	DynUNet	VesselFM	0.320±0.09*	0.322±0.08*	3.266±0.94*
		nnUNet	TopCoW	0.123±0.04*	0.079±0.02*	23.36±2.85*
		nnUNet	VesselV.	0.506±0.10	0.540±0.10*	2.757±1.62*
		nnUNet	DynaV.	**0.650**±0.08	**0.682**±0.09	**0.704**±0.35
TopBr.	Artery	DynUNet	VesselFM	0.374±0.12*	0.442±0.12*	2.470±0.94*
		nnUNet	TopCoW	0.192±0.04*	0.158±0.03*	21.15±1.92*
		nnUNet	VesselV.	0.455±0.12*	0.591±0.17*	8.070±4.29*
		nnUNet	DynaV.	**0.846**±0.05	**0.877**±0.06	**0.304**±0.15
	Vein	DynUNet	VesselFM	0.083±0.07*	0.210±0.17*	7.675±2.49*
		nnUNet	VesselV.	0.031±0.04*	0.132±0.15*	17.30±7.44*
		nnUNet	DynaV.	**0.957**±0.03	**0.974**±0.02	**0.08**±0.06
DynaV.	Artery	DynUNet	VesselFM	0.226±0.19*	0.284±0.22*	3.751±2.57*
		nnUNet	TopCoW	0.072±0.05*	0.054±0.04*	40.42±40.4*
		nnUNet	VesselV.	0.207±0.15*	0.276±0.21*	17.71±19.2*
		nnUNet	DynaV.	**0.757**±0.17	**0.510**±0.16	**2.28**±2.02
	Vein	DynUNet	VesselFM	0.064±0.05*	0.129±0.09*	5.137±1.25*
		nnUNet	VesselV.	0.028±0.03*	0.071±0.06*	29.12±19.9*
		nnUNet	DynaV.	**0.788**±0.09	**0.683**±0.12	**1.059**±0.8

Performance is assessed in two ways, namely: a) the binary vessel segmentation setting, where classes in the GTs and the predictions are merged into 0 = background, 1 = vessel; b) the artery vs vein setting, where classes in the GTs are kept as 0 = background, 1 = artery, 2 = vein.

3 Results

We compare DynaVessel to two training partitions of 112 images each from publicly available datasets: TopCoW [24] and VesselVerse [1]. We thus train

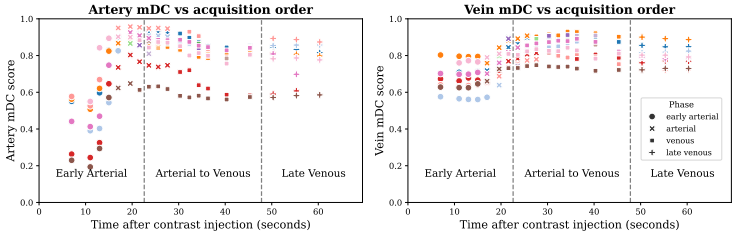

Fig. 4. mDC scores across phases. Each point represents a scan from a patient's dynamic 4D-CTA, with color encoding the patient. The transition between arterial and venous phases is not discrete, so we classify an scan's phase as arterial vs venous by taking the average HU intensity on the centerline of the GT artery/vein masks.

three nnUNet models as specified in Section 2.2 (from now on, we refer to the model trained on DynaVessel as M_{DV}). Likewise, we compare against the pretrained VesselFM [23] DynUNet, a foundation vessel segmentation model optimized on a collection of 928 images (including TopCoW) from vessel segmentation datasets encompassing various modalities, 10,000 images sampled from a generative model, and 500,000 randomly generated 3D images imitating blood vessels based on graphs derived from corrosion casts. Similar to VesselVerse, the labels in VesselFM are binary. We evaluate all models on five test datasets: 13 held-out samples from TopCoW [24]; 26 ISLES 2024 [16] scans annotated by the organizers of TopCoW; [24]; TopBrain's 26 scans; 13 held-out samples from VesselVerse [1]; and 14 held-out samples from DynaVessel. For DynaVessel, we use all available phases (165 images) to measure robustness to contrast variability.

In Fig. 3, we provide a qualitative overview of our results on a scan from TopBrain. M_{DV} achieves almost complete coverage of vascular structures, with few FNs. The difference in performance is larger on veins, which the other models mostly fail to detect. Although many FP detections are visible in Fig. 3(b, c, e, f), most are correctly segmented arterial and venous structures that are not annotated in TopBrain. Conversely, Fig. 3(a, d) shows that VesselFM generates many spurious segmentations, likely due to the limited amount of CTA it was trained on. Notably, TopBrain was acquired in Europe with Siemens scanners, while our dataset was acquired in the USA using a Canon-Toshiba scanner, highlighting the out-of-domain generalization potential of our approach.

In Table 1, we present quantitative results for vessel, artery, and vein segmentation. Despite being trained on a similarly-sized or smaller dataset, M_{DV} achieves higher performance on arteries and much higher performance on veins. Remarkably, training on DynaVessel leads to better results on data that is *in-domain* for the other models. Since the DynaVessel test set contains pre-contrast and early arterial phase data (see Fig. 4), all models performed worse on it compared to the single-phase TopBrain dataset. Lastly, in Fig. 5, we measure the impact of removing the various contrast phases used for training M_{DV} by optimizing two new models: M_{art} and M_{ven}, trained on 29 peak arterial and venous

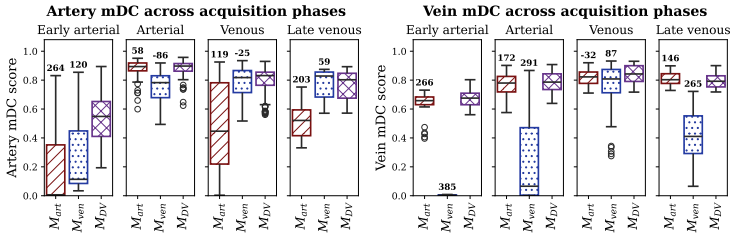

Fig. 5. mDC across contrast phases for M_{DV} and models trained either on the peak arterial (M_{art}) or peak venous (M_{ven}) phase images. The number above each boxplot codifies the delta between the mean Hounsfield unit (HU) intensity in arteries/veins during training vs the phase on which the model is evaluated. Single-phase training mDC correlates negatively with the delta between mean training and test HU.

phase images respectively. Crucially, M_{art} and M_{ven} have access to the same number of patients and GT annotations as M_{DV}. This test reveals that: 1) training on a single phase per patient reduces the models' capability to generalize across unseen phases; and 2) the model trained on all available phases often surpasses the other models even on the single phases used for training them.

4 Discussion

We proposed a semi-automatic data pre-processing method that utilizes dynamic CTA images acquired at multiple time points. Because this approach eliminates the need to manually annotate each time point, it significantly reduces the annotation burden and enhances training data diversity at no additional cost. Using this protocol, we created DynaVessel, a CTA intracranial artery and vein segmentation dataset comprising 275 images from 43 patients, covering a wide range of acquisition contrasts. We evaluate the impact of training a nnUNet model on our data, emphasizing sensitivity across vascular territories, and provide detailed insights on the model's ability to segment vessels across multiple phases.

The M_{DV} performance trends across dynamic CTA phases highlight the dependence of deep learning–based vessel segmentation on the phase of contrast enhancement, which is one of the principal sources of variation between real-world clinical CTA scans. When training on single-phase versions of DynaVessel, arterial segmentation performance is degraded by decreased conspicuity of arteries on late venous phase CTA scans, and venous segmentation is degraded by decreased conspicuity of veins in early arterial CTA scans. The difference in HU distributions between the training data and phase-specific test data correlates with performance; for instance, M_{ven} is trained on images with a mean venous attenuation of 485 HU, 87 HU higher than in the venous phase images ($mDC = 0.81$) and 385 HU greater than the mean venous HU in the early arterial phase ($mDC = 0.01$) images. These findings suggest that models trained on single-phase data rely on vessel attenuation related to phase of acquisition.

M_{DV} also achieved superior performance compared to models trained on datasets of similar or much larger sizes. The gap was more pronounced in veins, where M_{DV} achieved a mDC of 0.75 to 0.95, compared to under 0.01 for the other models. Nevertheless, M_{DV} also outperformed on arteries, with mDC scores of 0.75 to 0.86 vs scores ranging from 0.07 to 0.45 for the other models. The tSens and adHD metrics are consistent with this, showing M_{DV} beating the other models by a significant margin. Overall, we demonstrate a novel data-driven approach that provides a significant advance toward reliable automatic segmentation of all brain vascular structures in CTA, and will provide access to the code and training data required to reproduce our results. This advance has broad applicability as a primary or auxiliary input for AI pipelines designed to assess cerebrovascular pathologies. Perhaps more importantly, our work suggests that variation in phase of contrast enhancement may be a challenge to the robustness of deep learning models for assisted neurovascular CTA interpretation in clinical deployment, and points to a strategy for overcoming this challenge.

Acknowledgments. The authors acknowledge the financial support provided by NIH grant 1R01LM013891-01A1. Ceballos Arroyo, A. is grateful for the funding provided by Colombia's Minciencias and Fulbright. This work was partly supported by the DeltaAI system at the National Center for Supercomputing Applications through allocation CIS260200 from the ACCESS program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Falcetta, D., Marciano, V., Yang, K., *et al.* VesselVerse: A Dataset and Collaborative Framework for Vessel Annotation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. pp. 655–665. Springer Nature (2026). https://doi.org/10.1007/978-3-032-05169-1_63
2. Aydin, O.U., Taha, A.A., Hilbert, A., Khalil, A.A., Galinovic, I., Fiebach, J.B., Frey, D., Madai, V.I.: On the usage of average Hausdorff distance for segmentation performance assessment: hidden error when used for ranking. European Radiology Experimental **5**(1), 4 (Jan 2021), https://doi.org/10.1186/s41747-020-00200-2
3. Berger, A.H., Lux, L., Stucki, N., Bürgin, V., Shit, S., Banaszak, A., Rueckert, D., Bauer, U., Paetzold, J.C.: Topologically faithful multi-class segmentation in medical images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, 2024. pp. 721–731. Springer Nature (2024)
4. Ceballos-Arroyo, A.M., Nguyen, H.T., Zhu, F., Yadav, S.M., Kim, J., Qin, L., Young, G., Jiang, H.: Vessel-aware aneurysm detection using multi-scale deformable 3D attention. International Conference on Medical Image Computing and Computer-Assisted Intervention, 2024 **15005**, 754–765 (Oct 2024), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11986933/

5. Chen, L., Shaw, D.W.W., Dager, S.R., Corrigan, N.M., Chu, B., Kleinhans, N.M., Kuhl, P.K., Hwang, J.N., Yuan, C.: Quantitative assessment of the intracranial vasculature of infants and adults using iCafe (intracranial artery feature extraction). Frontiers in Neurology **Volume 12 - 2021** (2021), https://www.frontiersin.org/journals/neurology/articles/10.3389/fneur.2021.668298
6. Dundamadappa, S., Iyer, K., Agrawal, A., Choi, D.: Multiphase CT Angiography: A Useful Technique in Acute Stroke Imaging—Collaterals and Beyond. AJNR: American Journal of Neuroradiology **42**(2), 221–227 (Jan 2021), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7872177/
7. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature Methods **18**(2), 203–211 (Feb 2021), https://www.nature.com/articles/s41592-020-01008-z
8. Isensee, F., Wald, T., Ulrich, C., Baumgartner, M., Roy, S., Maier-Hein, K., Jäger, P.F.: nnU-Net revisited: A call for rigorous validation in 3d medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. p. 488–498. Springer-Verlag, Berlin, Heidelberg (2024), https://doi.org/10.1007/978-3-031-72114-4_47
9. Kapur, J.N., Sahoo, P.K., Wong, A.K.C.: A new method for gray-level picture thresholding using the entropy of the histogram. Computer Vision, Graphics, and Image Processing **29**(3), 273–285 (Mar 1985), https://www.sciencedirect.com/science/article/pii/0734189X85901252
10. Kassam, J., Thamm, F., Rist, L., Taubmann, O., Maier, A.: Detecting Large Vessel Occlusions using Graph Deep Learning. In: Proceedings of the First International Workshop on Geometric Deep Learning in Medical Image Analysis. pp. 149–159. PMLR (Nov 2022)
11. Kikinis, R., Pieper, S.D., Vosburgh, K.G.: 3D Slicer: A Platform for Subject-Specific Image Analysis, Visualization, and Clinical Support, pp. 277–289. Springer New York, New York, NY (2014), https://doi.org/10.1007/978-1-4614-7657-3_19
12. Kim, J., Ceballos-Arroyo, A., Lin, C.H., Liu, P., Jiang, H., Yadav, S., Wan, Q., Qin, L., Young, G.S.: Automated, anatomy-based, heuristic post-processing reduces false positives and improves interpretability of deep learning intracranial aneurysm detection models **16**(1), 3218, https://www.nature.com/articles/s41598-025-33083-7
13. Lal-Trehan Estrada, U.M., Oliver, A., Sheth, S.A., Lladó, X., Giancardo, L.: Strategies to combine 3D vasculature and brain CTA with deep neural networks: Application to LVO. iScience **27**(2), 108881 (Jan 2024). https://doi.org/10.1016/j.isci.2024.108881
14. Phansalkar, N., More, S., Sabale, A., Joshi, M.: Adaptive local thresholding for detection of nuclei in diversity stained cytology images. In: 2011 International Conference on Communications and Signal Processing. pp. 218–220 (Feb 2011), https://ieeexplore.ieee.org/document/5739305
15. Rajiah, P.S., Weber, N., Loewen, J., Kasten, H., Williamson, E., Moore, A., Leng, S.: Dynamic CT Angiography in Vascular Imaging: Principles and Applications. RadioGraphics **42**(7), E224–E225 (Nov 2022), https://pubs.rsna.org/doi/full/10.1148/rg.210177, publisher: Radiological Society of North America
16. Riedel, E.O., de la Rosa, E., Baran, T.A., *et al.*: ISLES'24 – a real-world longitudinal multimodal stroke dataset (2025), https://arxiv.org/abs/2408.11142

17. Shi, P., Hu, J., Yang, Y., Gao, Z., Liu, W., Ma, T.: Centerline boundary dice loss for vascular segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. p. 46–56. Springer-Verlag, Berlin, Heidelberg (2024)
18. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, J.P.W., Bauer, U., Menze, B.H.: clDice - a Novel Topology-Preserving Loss Function for Tubular Structure Segmentation. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16555–16564. IEEE, Nashville, TN, USA (Jun 2021), https://ieeexplore.ieee.org/document/9578225/
19. Talou, G.D.M., Safaei, S., Hunter, P.J., Blanco, P.J.: Adaptive constrained constructive optimisation for complex vascularisation processes. Scientific Reports **11**(1), 6180 (Mar 2021), https://www.nature.com/articles/s41598-021-85434-9, publisher: Nature Publishing Group
20. Tustison, N.J., Cook, P.A., Holbrook, A.J., *et al.*: The ANTsX ecosystem for quantitative biological and medical imaging. Scientific Reports **11**(1), 9068 (Apr 2021), https://doi.org/10.1038/s41598-021-87564-6
21. Wallu scheck, S., Canalini, L., Strohm, H., *et al.*: MR-CT multi-atlas registration guided by fully automated brain structure segmentation with CNNs. International Journal of Computer Assisted Radiology and Surgery **18**(3), 483–491 (Mar 2023), https://doi.org/10.1007/s11548-022-02786-x
22. Wang, Y., Narayanaswamy, A., Tsai, C.L., Roysam, B.: A broadly applicable 3-D neuron tracing method based on open-curve snake. Neuroinformatics **9**(2-3), 193–217 (Sep 2011). https://doi.org/10.1007/s12021-011-9110-5
23. Wittmann, B., Wattenberg, Y., Amiranashvili, T., Shit, S., Menze, B.: vesselFM: A foundation model for universal 3d blood vessel segmentation. In: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20874–20884 (2025). https://doi.org/10.1109/CVPR52734.2025.01944
24. Yang, K., Musio, F., Ma, Y., *et al.*: Benchmarking the CoW with the TopCoW Challenge: Topology-Aware Anatomical Segmentation of the Circle of Willis for CTA and MRA. ArXiv p. arXiv:2312.17670v4 (Jul 2025)
25. Yang, K., Shi, P., Huang, H., *et al.*: TopBrain Segmentation Challenge for Whole Brain Vessel Anatomy. medRxiv (2026). https://doi.org/10.64898/2026.05.28.26354312