Structural Congruence Matters: Biological Pretraining for Vascular Graph Extraction

Alessandro Scavone[1,2], Lucia Borrego[3,4][0009−0009−6439−8162], Vincenzo Marcianò[1,5][0009−0000−2238−8040], Christian Mata[4][0000−0003−4768−5062], Beatriz Giraldo[4][0000−0002−9910−8577], Josep Munuera[3][0000−0001−5154−1796], and Maria A. Zuluaga[1,5][0000−0002−1147−766X]

[1] EURECOM, Sophia Antipolis, France
[2] Politecnico di Torino, Turin, Italy
[3] Hospital Sant Pau, Barcelona, Spain
[4] Universitat Politecnica de Catalunya, Barcelona, Spain
[5] School of Biomedical Engineering & Life Sciences, King's College London, UK
maria.zuluaga@eurecom.fr

Abstract. Vascular graph extraction converts medical images into structured representations of vessel connectivity, enabling quantitative analysis of branching topology. However, training image-to-graph models requires expert-annotated vascular graphs, creating a substantial annotation bottleneck. Recent work leverages large non-medical datasets for structural pretraining, yet the structural alignment between source and vascular domains remains underexplored. We introduce a nature-to-nature pretraining paradigm that uses botanical branching networks as a biologically aligned source domain for vascular graph extraction. Leveraging large-scale annotated grapevine skeleton graphs and geometry-preserving graph representations, our approach promotes relational priors consistent with vascular growth and connectivity. Across heterogeneous 2D and 3D benchmarks, botanical pretraining consistently outperforms infrastructure-based pretraining under identical training settings. In addition, biologically aligned supervision achieves comparable performance with substantially fewer annotated vascular volumes. Multi-domain training further enhances robustness to previously unseen vascular datasets. These findings demonstrate that structural congruence between source and target domains provides a stronger inductive bias for vascular graph extraction than scale alone. Code available at: https://github.com/erc-caravel/Vascular-Graph-Extraction.

Keywords: structural-priors · vessels · image-to-graph.

1 Introduction

Vascular graph extraction infers structured network representations from medical images, revealing the connectivity of hemodynamic systems. Unlike pixel-wise segmentation, graph representations enable quantitative analysis of branching

topology and flow organization, which are critical for studying vascular remodeling and disease progression [4]. Despite advances in medical image segmentation, extracting topologically faithful vascular graphs remains challenging. A key bottleneck is the *annotation tax*: manually tracing and validating vascular graphs requires extensive expert effort, resulting in a severe imbalance between available imaging data and high-quality graph annotations.

To alleviate this constraint, recent work has leveraged large-scale infrastructure datasets, such as urban road networks, as auxiliary supervision for vascular graph learning [3]. This strategy aligns with foundation models (FMs), in which large, diverse pretraining data provide transferable inductive biases. However, data scale alone does not ensure structural alignment with the target domain. General-purpose FMs often underperform in medical settings, motivating the development of specialized Medical FMs (MedFMs) [30]. A similar limitation arises when transferring from man-made infrastructure to biological vasculature. Urban networks are shaped by human planning constraints, whereas vascular systems emerge from biological growth processes governed by flow efficiency and space-filling constraints. This mismatch in generative structure limits the suitability of infrastructure-derived priors for vascular graph extraction.

Building on [24], who distinguish between priors capturing superficial regularities and those that reflect deeper generative structure, we argue that overcoming the *annotation tax* requires improving representational congruence between source and target domains. Botanical transport systems exhibit structural principles (hierarchical branching, flow optimization, and space-filling organization) that more closely mirror vascular networks [19,25]. Thus, we propose to pretrain graph extraction models on large-scale plant transport networks to provide a biologically consistent inductive bias for cerebrovascular graph extraction.

In this work, we introduce a nature-to-nature pretraining paradigm for vascular graph extraction, grounded in structural congruence between source and target domains. Rather than transferring from structurally mismatched infrastructure data, we pretrain on botanical skeleton graphs to encode biologically consistent branching priors aligned with vascular organization. This encourages the model to internalize hierarchical, flow-constrained growth patterns characteristic of living transport networks. In addition, we adopt a geometrically faithful graph representation that preserves local curvature during training. By aligning the model's inductive bias with the generative structure of vascular systems, our approach enables high-fidelity graph extraction with fewer annotated vascular training samples than models pretrained on infrastructure data or purely semantic medical pretraining. We further investigate how these biologically aligned priors influence transferability, sample efficiency, and robustness across vascular domains.

Our contributions are threefold: (1) biologically aligned structural pretraining using botanical transport networks; (2) geometry-preserving graph representation strategies, including border-aware edge clipping and degree-2 node retention; and (3) empirical validation of improved performance, sample efficiency, and zero-shot generalization across heterogeneous 2D and 3D vascular datasets.

2 Related Work

Vascular Graph Extraction. Traditional vascular extraction relies on multi-stage pipelines combining segmentation, skeletonization, and graph pruning [8], which are prone to error propagation in noisy clinical settings. More recently, end-to-end image-to-graph frameworks have been proposed to directly infer structured connectivity from images [1,2,3,22,29]. Transformer-based architectures such as RelationFormer [22] jointly predict nodes and relations in a single stage, improving global topology reasoning. However, these approaches require substantial annotated graph data, which remains scarce in medical imaging.

Cross-Domain Pretraining for Image-to-Graph Transformers. To mitigate annotation scarcity, Berger et al. [3] introduced cross-domain and cross-dimensional transfer strategies for image-to-graph transformers. They demonstrated that large annotated road networks can serve as auxiliary supervision for vascular graph extraction, and proposed regularized edge sampling and adversarial domain adaptation to improve transferability. While this work establishes the feasibility of structural pretraining, road networks are shaped by planning constraints that differ from biological growth processes. However, the role of structural congruence between source and target domains remains unexplored.

Foundation Models and Data Scaling in Medical Imaging. Foundation model paradigms advocate large-scale pretraining on heterogeneous datasets to improve robustness and zero-shot generalization [14,27]. However, increasing data scale does not guarantee that learned representations capture the generative structure of the target domain [24]. In vascular segmentation, models such as vesselFM aggregate diverse real and synthetic data to enhance cross-domain performance. For graph extraction, preserving global connectivity and topology is essential, and naively combining datasets does not ensure structurally aligned relational priors [20,3].

Structural Congruence in Biological Transport Networks. Botanical and vascular systems arise under shared physical constraints, including flow efficiency, hierarchical branching, and space-filling organization [19,25]. These shared generative principles suggest a deeper structural alignment between plant and vascular transport networks compared to anthropogenic infrastructure. To our knowledge, prior work has not systematically investigated whether biologically aligned structural priors improve image-to-graph transformer pretraining for vascular extraction.

3 Method

Given an input image or volume $I \in \mathbb{R}^{\prod_{j=1}^{d} s_j}$, defined on a d-dimensional spatial grid of size $\mathbf{s} = (s_1, \ldots, s_d)$ with $d \in \{2, 3\}$, our goal is to predict a structured vascular graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. For 2D inputs, $d = 2$ and $\mathbf{s} = (H, W)$; for 3D inputs, $d = 3$ and $\mathbf{s} = (H, W, Z)$. The node set $\mathcal{V} = \{v_i\}_{i=1}^{N}$ consists of discrete centerline samples of the vascular graph, including endpoints, branching junctions,

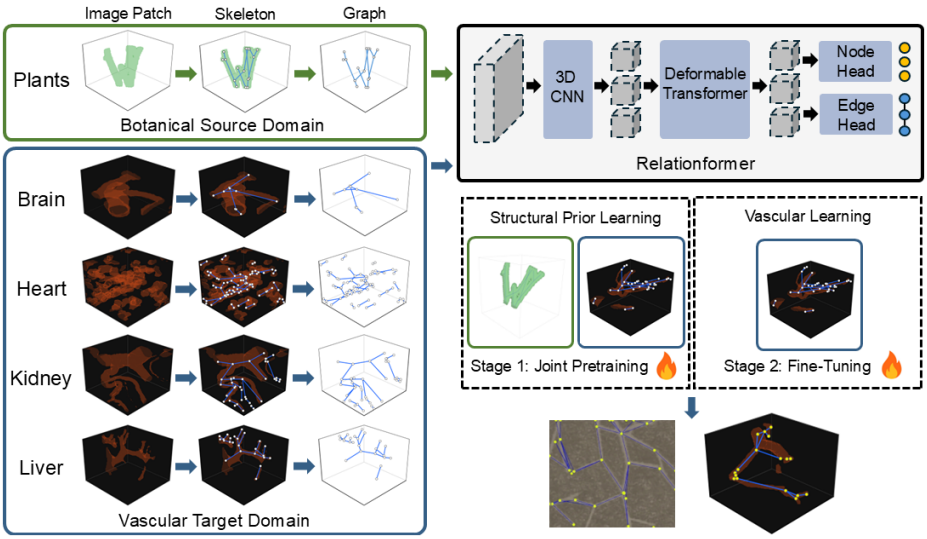

Fig. 1: Nature-to-nature pretraining paradigm. Botanical skeleton graphs provide biologically aligned structural supervision, whose learned relational priors are transferred to vascular graph extraction via fine-tuning.

and intermediate degree-2 nodes that preserve local geometric structure. Each node v_i is defined by normalized spatial coordinates $x_i \in [0,1]^d$ and a presence probability. The edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ encodes pairwise node connectivity, forming a topologically consistent vascular network.

3.1 Base Image-to-Graph Transformer

We build upon the RelationFormer architecture [22], a single-stage image-to-graph transformer that jointly predicts graph nodes and pairwise connectivity from image features. This avoids the multi-stage segmentation and skeletonization pipelines common in vascular graph extraction. It consists of a convolutional backbone followed by a transformer encoder-decoder. A fixed set of learnable object tokens predicts node locations, while a dedicated relation token captures global node interactions. The token-based formulation treats graph extraction as set prediction, enabling inference of a variable number of nodes and edges without hand-crafted post-processing.

For cross-domain and cross-dimensional transfer, we follow [3], which adapts RelationFormer to pretrain on structural source graphs before vascular fine-tuning. The framework augments standard node and edge prediction objectives with regularized edge sampling for connectivity imbalance and adversarial domain adaptation to improve source-target transfer. It provides a unified training interface for 2D/3D graph prediction, enabling structural supervision to generalize across domains. Within this setting, we vary the structural source domain and graph representation while keeping the transfer framework fixed.

3.2 Structural Modeling Choices

Our approach is motivated by the hypothesis that vascular graph extraction benefits from biologically aligned structural supervision and geometrically faithful graph representations. We therefore introduce modeling choices that reflect the generative properties of biological transport networks.

Structural Source Domain. Instead of pretraining on road networks [3], we use 25,900 annotated grapevine skeleton graph patches derived from 1,151 images [9]. Botanical transport systems exhibit hierarchical branching, flow-efficient organization, and space-filling structure consistent with vascular networks [15,25]. Pretraining on these graphs exposes the network to relational patterns governed by biological growth constraints rather than planning layouts.

Border-Aware Edge Clipping. When extracting fixed-size patches from graph annotations, edges intersecting patch boundaries are clipped at the intersection point by inserting a boundary node and reconnecting the truncated segment accordingly. This avoids artificial discontinuities introduced by naive edge removal and preserves consistent node-degree statistics within subgraphs.

Degree-2 Node Retention. Standard pipelines collapse degree-2 nodes into straight-line segments between junctions. While suitable for infrastructure networks, this removes curvature information characteristic of vascular anatomy. We retain degree-2 nodes during preprocessing to preserve local geometry and enable curvature-aware relational supervision.

3.3 Cross-Dimensional Rendering

Structural source graphs are defined in 2D, whereas several vascular targets are volumetric ($d = 3$). To enable pretraining of 3D backbones from 2D structural data, we adopt the 2D-to-3D transfer framework introduced by Berger et al. [3].

Each graph is rasterized into a thin skeleton mask M_{skel}. To approximate tubular vessel geometry and partial volume effects observed in volumetric imaging, the mask is iteratively convolved with a 3×3 all-ones kernel K_{ones}:

$$M_i = \text{Conv2d}(M_{i-1}, K_{\text{ones}}), \quad i \in \{1, \ldots, k\}, \tag{1}$$

$$M_{\text{thick}} = \text{clip}(M_k, 0, 255). \tag{2}$$

The resulting thickened mask provides a pseudo-volumetric representation compatible with 3D convolutional backbones that preserves branching topology.

3.4 Optimization and Training Setup

Training proceeds in two phases. First, structural source samples (plants) and vascular target samples are optimized jointly. Second, training continues using only vascular datasets, with all model parameters remaining trainable.

Models are trained on fixed-size patches interpreted as structured subgraphs. To ensure stable relational supervision, we retain only patches containing 6–70 nodes, avoiding overly sparse or excessively dense graph configurations.

4 Experiments and Results

We evaluate the impact of structurally aligned pretraining and graph representation refinements on vascular graph extraction across 2D and 3D datasets.

4.1 Experimental Setup

Datasets. We use two 2D structural source domains: 20-Cities [11] (180 images/99,200 road-graph patches) and PlantsGuyot3d2 [9] (1,511 images/25,900 grapevine skeleton patches). Target domains comprise both 2D and 3D vascular datasets across organs and modalities. The 2D target is OctaSynth [17]. The 3D targets are: TopCoW (CT/MRA) [28], TubeTK [6], IXI dataset[6], DeepVess [10], TubeNet [12], syntheticMRI [21], VesselExpress [23], and HRKidney [13].

Preprocessing and Graph Construction. All datasets are standardized into a unified relational format. Images are resampled to a common coordinate space and partitioned into fixed-size patches of 128^2 for 2D and 64^3 for 3D experiments. When graph annotations are not provided, ground-truth graphs are extracted from segmentation masks using the Voreen vessel graph extraction method [8,18], followed by processing with the unified graph representation described in Sec. 3.

Evaluation Metrics. We evaluate graph reconstruction using both graph similarity and detection metrics. For graph similarity, we report the TOPO score [5], which evaluates topological correctness and geometric fidelity in 2D, and the Street Mover Distance (SMD), which approximates the Wasserstein distance between embedded graphs to measure structural discrepancy. For node and edge detection, we report mean Average Precision (mAP) and mean Average Recall (mAR). Following prior image-to-graph work [3,22], , nodes are evaluated as fixed-size bounding boxes centered at predicted and ground-truth node locations, with side length $\Delta x = 0.2$ in normalized coordinates. Edges are represented by axis-aligned bounding boxes spanning their endpoint nodes. Node and edge matches are then determined using the corresponding IoU-based criterion.

Implementation Details. We use the RelationFormer architecture [22] with a 3D-SEResNet backbone [16], employing 80/2 object–relation tokens for 2D experiments (5 relation tokens for OctaSynth) and 120/2 for 3D. In both 2D and 3D experiments, we follow the optimization strategy outlined in [3] by adopting AdamW as the optimizer and batch size 32.

4.2 Results

Structural Source-Domain Comparison. We compare botanical and road networks as structural initialization sources under identical training conditions. Roads are evaluated both at their full scale (99.2k patches) and in a size-matched setting (25.9k patches). Border-aware clipping is applied uniformly across all

[6] https://brain-development.org/ixi-dataset/

Table 1: Comparison of plant and road structural priors on 2D (OctaSynth) and 3D (syntheticMRI) targets. Rows labeled "+2-degree" retain degree-2 nodes in the source graphs at preprocessing. Gray denotes models without pretraining.

Source	Num. Source Samples	SMD↓	Node mAP↑	Node mAR↑	Edge mAP↑	Edge mAR↑	TOPO Prec.↑	TOPO Rec.↑
2D Target: OctaSynth								
No pretrain	-	0.005	0.248	0.353	0.152	0.256	0.719	0.720
roads	99.2k	0.005	0.307	0.398	0.205	0.304	0.773	0.752
roads	25.9k	0.006	0.251	0.352	0.166	0.245	0.754	0.682
plants	25.9k	0.003	0.516	0.581	0.378	0.463	0.828	0.749
roads+2-degree	25.9k	0.002	0.476	0.552	0.369	0.472	0.792	0.830
plants+2-degree	25.9k	**0.0006**	**0.629**	**0.692**	**0.553**	**0.630**	**0.843**	**0.884**
3D Target: syntheticMRI								
No pretrain	-	0.021	0.194	0.301	0.017	0.059	-	-
roads	99.2k	**0.014**	0.406	0.489	0.183	0.257	-	-
roads	25.9k	0.021	0.260	0.322	0.091	0.141	-	-
plants	25.9k	0.017	0.396	0.469	0.167	0.239	-	-
roads+2-degree	25.9k	0.015	0.449	0.520	0.180	0.253	-	-
plants+2-degree	25.9k	0.015	**0.462**	**0.530**	**0.202**	**0.282**	-	-

source domains, and degree-2 node retention is evaluated symmetrically for both roads and plants.

Pretraining on either domain improves performance over training without structural supervision. However, botanical pretraining consistently outperforms road-based pretraining on both 2D (OctaSynth) and 3D (syntheticMRI) targets (Table 1), including in the size-matched setting. Moreover, plants+2-degree outperforms roads+2-degree under identical graph representation and training settings. Together, these results suggest that biologically aligned structural priors provide more transferable relational supervision than infrastructure-based ones within the studied framework.

Retaining degree-2 nodes further improves performance for both sources, including on syntheticMRI, whose target graphs do not retain degree-2 nodes, suggesting improved modeling of local connectivity patterns.

Sample Efficiency Assessment. We evaluate how structural pretraining influences label efficiency by varying the amount of annotated target data used for fine-tuning on IXI. Starting from a fixed pretrained initialization (roads or plants), we fine-tune the model using fractions $\{100\%, 75\%, 50\%, 30\%, 15\%\}$ of the IXI training set, subsampled at the volume level to preserve spatial coherence and graph statistics. Validation and test splits remain unchanged.

Figure 2 (left) reports performance on the held-out IXI test set. Plant-pretrained models consistently outperform road-pretrained models across all annotation budgets. Remarkably, a plant-pretrained model using just 30% data already approaches the performance of a road-pretrained model trained on the full

Table 2: Cumulative multi-domain training configurations used in zero-shot experiments. Each configuration increases anatomical and modality diversity.

Config.	Included Training Domains
E1	syntheticMRI, TopCoW-CT, TopCoW-MRA, HiP-CT, MiniVess
E2	E1 + VesselExpress-Bladder, VesselExpress-Heart, tubenet-MRI
E3	E2 + tubenet-2photon, DeepVess, 3Dircadb1, TubeTK
E4	E3 + VesselExpress-Brain, IXI (train-val), HRKidney

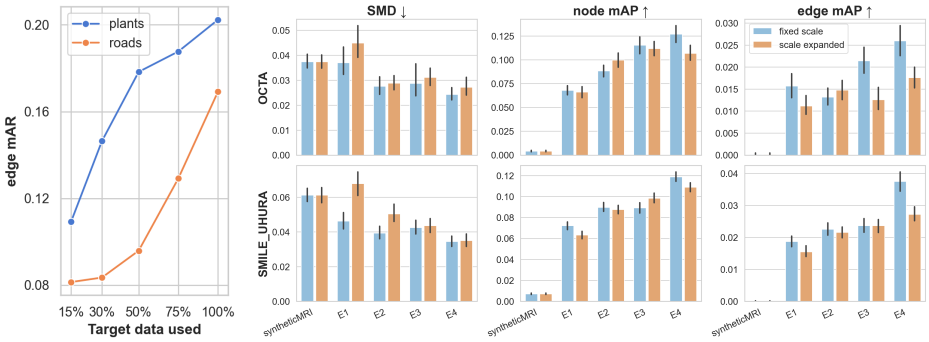

Fig. 2: Left: IXI performance under varying annotation budgets for road- and plant-pretrained models. Right: Zero-shot performance on unseen datasets across cumulative multi-domain training configurations (E1–E4).

dataset (100%). While the performance gap is tighter at the extreme boundaries (15% and 100%), it peaks significantly across intermediate budgets, proving that biologically aligned structural priors reduce the amount of vascular annotation required to achieve competitive graph-extraction performance.

Zero-Shot Generalization. We assess generalization to unseen vascular datasets without target-specific fine-tuning by disentangling the effects of domain diversity and training scale. We compare (i) single-domain training on ∼4k syntheticMRI patches; (ii) fixed-scale multi-domain training, where the total pool remains ∼4k patches but incorporates auxiliary datasets (up to 300 patches per domain); and (iii) scale-expanded multi-domain training, where additional domains are added without reducing the syntheticMRI subset. For (ii) and (iii), cumulative configurations ($E1$–$E4$) progressively increase anatomical and modality diversity (Table 2). Evaluation is performed on held-out SMILE_UHURA [7] and 3D OCTA [26] datasets.

As shown in Fig. 2 (right), increasing training diversity consistently improves zero-shot performance. Even at a fixed scale, introducing heterogeneous vascular domains yields substantial gains, while further improvements are observed when diversity is combined with increased data volume. These results indicate that diversity is the primary driver of cross-dataset robustness, with scale providing complementary benefits.

5 Conclusions

We demonstrated that structural congruence between source and target domains is a key factor in cross-domain pretraining for vascular graph extraction. Replacing infrastructure-based supervision with botanically aligned transport networks yields stronger and more sample-efficient graph prediction despite fewer source samples. Combined with geometrically faithful graph representations, biological pretraining improves relational learning across dimensions.

These findings suggest that effective pretraining for structured medical prediction should prioritize generative alignment, not only dataset scale. In vascular graph extraction, source domains governed by biological growth and transport principles provide more useful relational priors than visually convenient but structurally mismatched alternatives. Future work will evaluate whether these benefits extend to other image-to-graph architectures, skeletonization pipelines, and biological transport networks beyond grapevine skeletons.

Acknowledgments. This work is co-funded by the European Union (ERC, CARAVEL, 101171357) and by the European project SEARCH (No. 101172997), which is supported by the Innovative Health Initiative Joint Undertaking (IHI JU). We acknowledge computational resources provided by GENCI-IDRIS on Jean Zay A100 partition (grants 2025-AD011016543, 2026-A0200617550).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the EU or the European Research Council. Neither the EU nor the granting authority can be held responsible for them.

References

1. Almasi, S., Xu, X., Ben-Zvi, A., Lacoste, B., Gu, C., Miller, E.L.: A novel method for identifying a graph-based representation of 3-d microvascular networks from fluorescence microscopy image stacks. Medical image analysis **20**(1), 208–223 (2015)
2. Batten, J., Sinclair, M., Bai, Y., Glocker, B., Schaap, M.: Image-to-tree with recursive prompting. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2024)
3. Berger, A.H., Lux, L., Shit, S., Ezhov, I., Kaissis, G., Menten, M.J., Rueckert, D., Paetzold, J.C.: Cross-domain and cross-dimension learning for image-to-graph transformers. arXiv preprint arXiv:2403.06601 (2024)
4. Bernabeu, M.O., Jones, M.L., Nielsen, J.H., Krüger, T., Nash, R.W., Groen, D., Schmieschek, S., Hetherington, J., Gerhardt, H., Franco, C.A., et al.: Computer simulations reveal complex distribution of haemodynamic forces in a mouse retina model of angiogenesis. Journal of The Royal Society Interface **11**(99) (2014)
5. Biagioni, J., Eriksson, J.: Inferring road maps from global positioning system traces: Survey and comparative evaluation. Transportation research record **2291**(1), 61–71 (2012)
6. Bullitt, E., Zeng, D., Gerig, G., Aylward, S., Joshi, S., Smith, J.K., Lin, W., Ewend, M.G.: Vessel tortuosity and brain tumor malignancy: a blinded study1. Academic radiology **12**(10), 1232–1240 (2005)

7. Chatterjee, S., Mattern, H., Dörner, M., Sciarra, A., Dubost, F., Schnurre, H., Khatun, R., Yu, C.C., Hsieh, T.L., Tsai, Y.S., et al.: Smile-uhura challenge–small vessel segmentation at mesoscopic scale from ultra-high resolution 7t magnetic resonance angiograms. arXiv preprint arXiv:2411.09593 (2024)
8. Drees, D., Scherzinger, A., Hägerling, R., Kiefer, F., Jiang, X.: Scalable robust graph and feature extraction for arbitrary vessel networks in large volumetric datasets. BMC bioinformatics **22**(1), 346 (2021)
9. Gentilhomme, T., Villamizar, M., Corre, J., Odobez, J.M.: Towards smart pruning: Vinet, a deep-learning approach for grapevine structure estimation. Computers and Electronics in Agriculture **207**, 107736 (2023)
10. Haft-Javaherian, M., Fang, L., Muse, V., Schaffer, C.B., Nishimura, N., Sabuncu, M.R.: Deep convolutional neural networks for segmenting 3d in vivo multiphoton images of vasculature in alzheimer disease mouse models. PloS one **14**(3), e0213539 (2019)
11. He, S., Bastani, F., Jagwani, S., Alizadeh, M., Balakrishnan, H., Chawla, S., Elshrif, M.M., Madden, S., Sadeghi, M.A.: Sat2graph: Road graph extraction through graph-tensor encoding. In: European Conference on Computer Vision. pp. 51–67. Springer (2020)
12. Holroyd, N.A., Li, Z., Walsh, C., Brown, E., Shipley, R.J., Walker-Samuel, S.: tubenet: a generalizable deep learning tool for 3d vessel segmentation. Biology Methods and Protocols **10**(1), bpaf087 (2025)
13. Kuo, W., Rossinelli, D., Schulz, G., Wenger, R.H., Hieber, S., Müller, B., Kurtcuoglu, V.: Terabyte-scale supervised 3d training and benchmarking dataset of the mouse kidney. Scientific data **10**(1), 510 (2023)
14. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. Nature Communications **15**(1), 654 (2024)
15. McCulloh, K.A., Sperry, J.S.: Murray's law and the vascular architecture of plants. Ecology and biomechanics. Boca Raton, FL, USA: Taylor and Francis pp. 105–120 (2006)
16. Meng, L., Zhou, L., Zhang, W., Xie, K.: Robust epileptic seizure prediction: A 3d-seresnet framework for patient-specific and multi-patient generalization. iScience **28**(12) (2025)
17. Menten, M.J., Paetzold, J.C., Dima, A., Menze, B.H., Knier, B., Rueckert, D.: Physiology-based simulation of the retinal vasculature enables annotation-free segmentation of oct angiographs. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 330–340. Springer (2022)
18. Meyer-Spradow, J., Ropinski, T., Mensmann, J., Hinrichs, K.: Voreen: A rapid-prototyping environment for ray-casting-based volume visualizations. IEEE Computer Graphics and Applications **29**(6), 6–13 (2009)
19. Murray, C.D.: The physiological principle of minimum work: I. the vascular system and the cost of blood volume. Proceedings of the National Academy of Sciences **12**(3), 207–214 (1926)
20. Paetzold, J.C., McGinnis, J., Shit, S., Ezhov, I., Büschl, P., Prabhakar, C., Sekuboyina, A., Todorov, M., Kaissis, G., Ertürk, A., et al.: Whole brain vessel graphs: A dataset and benchmark for graph learning and neuroscience. In: Thirty-fifth conference on neural information processing systems datasets and benchmarks track (Round 2) (2021)
21. Schneider, M., Reichold, J., Weber, B., Székely, G., Hirsch, S.: Tissue metabolism driven arterial tree generation. Medical image analysis **16**(7), 1397–1414 (2012)

22. Shit, S., Koner, R., Wittmann, B., Paetzold, J., Ezhov, I., Li, H., Pan, J., Sharifzadeh, S., Kaissis, G., Tresp, V., Menze, B.: Relationformer: A unified framework for image-to-graph generation. arXiv preprint arXiv:2203.10202 (2022)
23. Spangenberg, P., Hagemann, N., Squire, A., Förster, N., Krauß, S.D., Qi, Y., Yusuf, A.M., Wang, J., Grüneboom, A., Kowitz, L., et al.: Rapid and fully automated blood vasculature analysis in 3d light-sheet image volumes of different organs. Cell Reports Methods **3**(3) (2023)
24. Vafa, K., Chang, P.G., Rambachan, A., Mullainathan, S.: What has a foundation model found? inductive bias reveals world models. In: Forty-second International Conference on Machine Learning (2025), https://openreview.net/forum?id=i9npQatSev
25. West, G.B., Brown, J.H., Enquist, B.J.: A general model for the origin of allometric scaling laws in biology. Science **276**(5309), 122–126 (1997)
26. Wittmann, B., Glandorf, L., Paetzold, J.C., Amiranashvili, T., Wälchli, T., Razansky, D., Menze, B.: Simulation-based segmentation of blood vessels in cerebral 3d octa images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 645–655. Springer (2024)
27. Wittmann, B., Wattenberg, Y., Amiranashvili, T., Shit, S., Menze, B.: vesselfm: A foundation model for universal 3d blood vessel segmentation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 20874–20884 (2025)
28. Yang, K., Musio, F., Ma, Y., Juchler, N., Paetzold, J.C., Al-Maskari, R., Höher, L., Li, H.B., Hamamci, I.E., Sekuboyina, A., et al.: Benchmarking the cow with the topcow challenge: Topology-aware anatomical segmentation of the circle of willis for cta and mra. ArXiv pp. arXiv–2312 (2025)
29. Yao, L., Shi, F., Wang, S., Zhang, X., Xue, Z., Cao, X., Zhan, Y., Chen, L., Chen, Y., Song, B., Wang, Q., Shen, D.: Tag-net: Topology-aware graph network for centerline-based vessel labeling. IEEE Transactions on Medical Imaging **42**(11), 3155–3166 (2023)
30. Zhang, S., Metaxas, D.: On the challenges and perspectives of foundation models for medical image analysis. Medical image analysis **91**, 102996 (2024)