Trexplorer Super: Topologically Correct Centerline Tree Tracking of Tubular Objects in CT Volumes

Roman Naeem[0009−0001−2625−125X], David Hagerman[0000−0002−5193−3789], Jennifer Alvén[0000−0003−4195−9325], Lennart Svensson[0000−0003−0206−9186], and Fredrik Kahl[0000−0001−9835−3020]

Chalmers University of Technology, 412 96 Gothenburg, Sweden
{nroman,david.hagerman,alven,lennart.svensson,fredrik.kahl}@chalmers.se

Abstract. Tubular tree structures, such as blood vessels and airways, are essential in human anatomy, and accurately tracking them while preserving their topology is crucial for various downstream tasks. Trexplorer is a recurrent model designed for centerline tracking in 3D medical images, but it is prone to predicting duplicate branches and terminating tracking prematurely. To address these issues, we present Trexplorer Super, an enhanced version that substantially improves performance through several novel advancements. Evaluating centerline tracking models is challenging due to the lack of public benchmark datasets. To enable thorough evaluation, we develop three centerline datasets, one synthetic and two real, each with increasing difficulty. Using these datasets, we perform a comprehensive comparison of existing state-of-the-art (SOTA) models with our approach. Trexplorer Super outperforms previous SOTA models on every dataset. Our results also highlight that strong performance on synthetic data does not necessarily translate to real datasets. The code and datasets are available at
https://github.com/RomStriker/Trexplorer-Super

Keywords: centerline tracking · tubular structures · tree topology.

1 Introduction

Tubular tree structures in the vascular and respiratory systems play a critical role in transporting essential substances throughout the body. Accurately tracking the centerlines of these structures in medical images is fundamental for early diagnosis, treatment, and various downstream tasks [9,5,6,2]. In this paper, we introduce a new method for centerline tree tracking and propose a comprehensive framework for its evaluation.

Several existing approaches address the challenge of centerline extraction, but each comes with its own limitations. A common method segments the image and then applies skeletonization [18], but such models struggle with long-range dependencies, leading to connectivity issues. Other models [11,16] detect centerline nodes and edges in a two-step process but also suffer from connectivity

errors. Recurrent models, such as reinforcement learning-based methods [23,7], iteratively track centerlines but rely on complex pipelines. Trexplorer [10] simplifies this with a DETR-based transformer [1] that uses breadth-first tracking to ensure correct topology. However, it struggles with duplicate branch detections and premature tracking terminations.

To overcome the limitations of existing centerline tracking methods, we propose Trexplorer Super, which builds on the Trexplorer model with several key enhancements to improve accuracy, robustness, and completeness. Our method reduces premature terminations and duplicate branches while improving new branch detection and preserving fine spatial details in image features. To ensure more consistent centerline extraction, we introduce Super Trajectory Training, a strategy that retains and reuses tracking information across multiple steps. We also refine feature representation with Focal Cross Attention, which selectively attends to a focal region in high-resolution image features while maintaining broader contextual awareness. To further enhance robustness, we employ Target Augmentation, a strategy that improves bifurcation and new branch detection while minimizing duplicate branches. These advancements contribute to a more reliable and comprehensive centerline tracking framework.

Evaluating centerline tracking in 3D medical images is challenging due to the lack of publicly available real datasets. Existing synthetic datasets have topological limitations, and strong performance on these does not generalize well to real data. To address this, we create one synthetic and two real datasets and establish a comprehensive baseline by evaluating prior SOTA models alongside our approach using point-, branch-, and tree-level metrics.

Our key contributions include: **(1)** enhancing the Trexplorer framework with novel techniques, namely Super Trajectory Training, Focal Cross Attention, and Target Augmentation; and **(2)** creating three datasets and thoroughly evaluating the previous SOTA models alongside our method.

2 Method

Our goal is to estimate the centerline tree from a given CT volume and a starting root point. The centerline tree is represented as a graph (V, E) with V nodes and E edges. Each node $\mathbf{v} \in V$ is defined as a vector $\mathbf{v} = [x, y, z, r]$, representing the 3D position and radius of a centerline point, while an edge $\mathbf{e} \in E$ represents a connection between two nodes.

Trexplorer Super is a DETR-based model that uses object queries to track branches. It begins at the root point and tracks each branch to its endpoint. The model estimates the total number of branches and the number of nodes within each branch. Tracking follows a sequential, breadth-first approach in which, at each step, all child nodes at the next graph level are predicted. Each child node is approximately one voxel away from its parent and is classified into one of three categories: end node, intermediate node, or bifurcation node. The model stops tracking a branch when it predicts an end node. If the node is intermediate, it continues tracking subsequent nodes. For bifurcation nodes, it halts tracking of

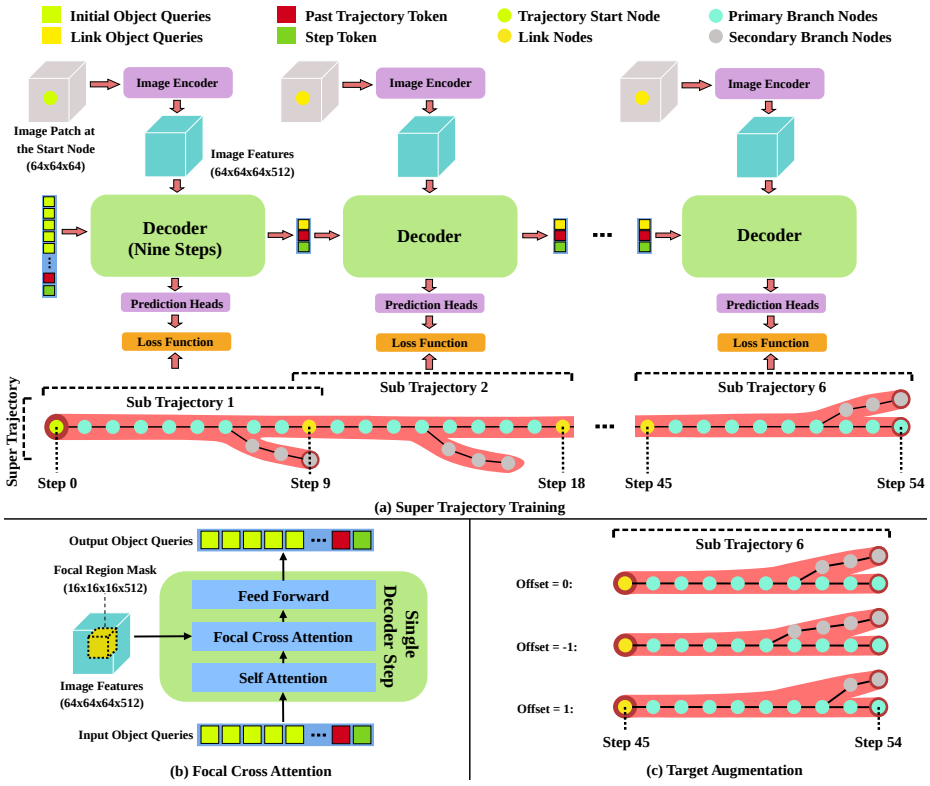

Fig. 1. Overview of the Trexplorer Super architecture, illustrating: (a) Super Trajectory Training, (b) Focal Cross Attention, and (c) Target Augmentation.

the current branch and assigns a new set of 26 object queries, some of which start tracking the new branches, while the rest are discarded. Trexplorer Super builds on Trexplorer by incorporating the following novel key components. The model architecture is shown in Figure 1.

2.1 Super Trajectory Training

Trexplorer generates a centerline graph over 9 steps, starting from the center of a volume patch. New patches are created at the last step nodes of the tracked graph, which are used by the model, along with the past trajectory, to continue tracking until all endpoints of the tree are reached. The past trajectory token, an embedding of up to 10 previous node positions, is the only source of past information when tracking in a new patch. This simplifies the formation of training batches, as each batch is independent. However, it results in the loss of the learned past trajectory embedding stored in the object queries, causing premature branch terminations and missed branches, leading to a lower recall.

To better utilize past trajectory information, we propose Super Trajectory Training (STT). In STT, each sample consists of a super trajectory of size 54, divided into six sub-trajectories of size 10 with shared link nodes. Each sub-trajectory is paired with a volume patch centered at its starting node, as shown in Figure 1(a). Secondary branches may appear within a sub-trajectory, but they are only tracked within the corresponding patch during training.

During training, object query outputs from the previous sub-trajectory are used as inputs for the next sub-trajectory, along with new volume patches. This preserves the valuable past trajectory information embedded in the object queries. In contrast, Trexplorer's training strategy is analogous to training on a single sub-trajectory. The first sub-trajectory still relies on the past trajectory token. During inference, tracking begins at the root point, and for subsequent patches, the object query outputs from previous patches are used. This forms a continuous chain from the root to all endpoints, enabling more effective tracking.

2.2 Focal Cross Attention

In Trexplorer, object queries track branches by aggregating information from image features using the transformer's cross-attention module [19]. For tubular structures like vessel trees, capturing long-range dependencies and fine-grained spatial details is crucial for accurately locating thin, elongated branches. However, using large, high-resolution features quickly becomes computationally infeasible for 3D medical images. Some methods address this by using learned sparse attention [26,13,24], but they rely on object queries to determine which features to attend to. This conflicts with Trexplorer's use of object queries to store branch tracking history, leading to poor performance.

Trexplorer Super introduces Focal Cross Attention (FCA), which extracts high-resolution features over a large region but restricts cross-attention to the small focal region where branches are being tracked, as shown in Figure 1(b). By training the model end-to-end, the responsibility for aggregating long-range dependencies while maintaining fine-grained spatial details is delegated to the feature extractor. This design allows the decoder object queries to focus on retaining tracking history while cross-attending to a smaller, more relevant set of features.

2.3 Target Augmentation

As the radius of a bifurcation node increases, the area of viable positions for the bifurcation also increases, rather than being a single fixed point. Trexplorer Super is trained to account for this positional ambiguity using Target Augmentation. During target augmentation, for each bifurcation point in the primary branch of a super trajectory, an offset is sampled from a Laplace distribution with mean $\mu = 0$ and scale b proportional to the bifurcation radius. This offset shifts secondary branches up or down along the primary branch, effectively generating viable augmented targets as shown in Figure 1(c). The branches around these new bifurcation points are smoothed to preserve natural trajectories.

Table 1. Statistics of the ground truth centerline graphs for the provided datasets.

Dataset	Samples			Max Node	Average			Radius		
	train	val	test	Degree	points	depth	width	max	mean	min
Synthetic Dataset	368	32	100	4	2 763	343	20	17	5.37	2
ATM'22	220	16	60	4	7 398	504	55	23	2.47	1
Parse 2022	72	8	20	4	19 644	498	126	38	2.55	1

Training with augmented targets encourages the model to consider more nodes as potential bifurcations, improving the detection of new branches. This strategy also reduces duplicates, leading to fewer volume patches to process and eliminating the need for post-processing. The reduction in duplicates can be attributed to the combined effect of Super Trajectory Training and Target Augmentation, which greatly enhances the effect of the Hungarian loss and self-attention between object queries, two key components used for preventing duplicates in Trexplorer. Additionally, a small amount of Gaussian noise ($\mu = 0, \sigma = 0.025$) is added to all points to further improve robustness.

3 Experiments and Results

3.1 Datasets

To our knowledge, the only publicly available 3D centerline tree dataset is the synthetic vessel tree dataset [18]. However, these trees were generated without collision avoidance, resulting in unrealistic vessel intersections. To address this issue, we use the Synthetic Vascular Toolkit (SVT) [15,14] to generate a new synthetic tree dataset with collision avoidance. We follow the same procedure as [18] to create the corresponding images. This dataset serves as a useful toy example for model research and development.

For a comprehensive evaluation, we also generate centerline ground truth from two publicly available real tubular tree segmentation datasets: the ATM'22 dataset [21,22,25,20,12] (airway segmentation), licensed under CC BY-NC and the Parse 2022 dataset [8] (pulmonary artery segmentation), licensed under CC BY-NC-ND 6.0. The use of these datasets complies with the terms set by the dataset owners. Before extracting the ground truth, we resample the data volumes to 0.5 mm isotropic resolution. The Vascular Modeling Toolkit (VMTK) [4,3] is used to extract root points from segmentation masks, which Kimimaro [17] then uses to trace the tubular tree centerlines. Further dataset statistics are provided in Table 1.

3.2 Evaluation Metrics

We evaluate predictions at the point, branch, and tree levels to capture different aspects of accuracy. At the point level, we use Precision, Recall, and F1-score. A predicted node is considered a True Positive (TP) if a ground truth node

exists within its 1.5-voxel radius that has not been matched to another prediction; otherwise, it is a False Positive (FP). An unmatched ground truth node is considered a False Negative (FN). We also assess radius accuracy using Mean Absolute Error (MAE). Given the large number of nodes, we avoid metrics that require solving the Linear Assignment Problem.

At the branch level, we again use Precision, Recall, and F1-score while treating branches as objects. A predicted branch is considered a TP if it correctly matches at least 80% of the points in a ground truth branch within a 1.5-voxels radius, provided that the ground truth branch has not already been matched. Otherwise, it is an FP. Unmatched ground truth branches are considered FNs. To evaluate the overall graph structure, we use topological metrics, specifically, the MAE of Betti-0 (connected components) and Betti-1 (cycles).

3.3 Experiments

We evaluate two previous state-of-the-art (SOTA) models, Vesselformer and Trexplorer, alongside our proposed method, Trexplorer Super, on three datasets. All models were trained on a single node with four A100 GPUs, with Trexplorer and Trexplorer Super using mixed precision for improved efficiency. To ensure a fair comparison focused on model improvements, Trexplorer and Trexplorer Super share nearly identical hyperparameters: allocating 26 tokens for a bifurcation node, and a maximum of 196 tokens. Both models were trained for approximately 2 million iterations. Vesselformer produced the best results for author-optimized hyperparameters, with 80 object tokens and approximately 12 million training iterations. Each model was trained five times per dataset, and we report the mean and standard deviation for each metric.

3.4 Results

Tables 2 and 3 summarize the performance of the evaluated models on the proposed datasets, while Figure 2 provides a visual comparison using one sample from each dataset. On the synthetic dataset, Vesselformer achieves a higher F1-score than Trexplorer. Although Trexplorer has better recall, it generates many duplicate branches, leading to low precision. Trexplorer Super improves recall over Vesselformer, though still lower than Trexplorer, while drastically reducing duplicates, resulting in the highest overall F1-score. On ATM'22, Vesselformer retains some centerline tracking ability but misses branches and predicts duplicates, resulting in a low overall score. Trexplorer performs worse, barely tracking any branches due to limited access to past trajectory information. Trexplorer Super, with its enhancements, delivers strong improvements across all metrics. Parse 2022 presents a greater challenge due to denser trees and weaker vessel signals. Both Vesselformer and Trexplorer struggle, but Trexplorer Super clearly outperforms them.

Trexplorer Super achieves the lowest radius MAE for all datasets. At branch level, it also performs the best across all metrics. Both Trexplorer and Trexplorer Super ensure topological correctness with zero Betti-0 and Betti-1 errors,

Table 2. Comparison of different models using point-level metrics on the Synthetic, ATM'22, and Parse 2022 datasets.

Model	Point-level			
	Precision(%)↑	Recall(%)↑	F1(%)↑	Radius (MAE)↓
Synthetic Dataset				
Vesselformer	44.53 ± 7.87	61.52 ± 1.14	48.18 ± 5.62	0.4244 ± 0.0134
Trexplorer	30.91 ± 9.45	**78.21 ± 4.13**	39.40 ± 8.62	0.2263 ± 0.0323
Trexplorer Super	**91.91 ± 3.28**	70.44 ± 3.02	**77.83 ± 1.89**	**0.0955 ± 0.0061**
ATM'22 Dataset				
Vesselformer	22.32 ± 1.35	34.37 ± 0.94	26.77 ± 1.27	0.7908 ± 0.0095
Trexplorer	3.20 ± 0.73	4.33 ± 0.63	3.34 ± 0.30	0.9744 ± 0.0844
Trexplorer Super	**67.51 ± 1.35**	**60.65 ± 2.01**	**60.45 ± 1.03**	**0.3925 ± 0.0241**
Parse 2022 Dataset				
Vesselformer	18.49 ± 1.84	15.28 ± 0.83	16.43 ± 0.78	1.1144 ± 0.0269
Trexplorer	9.87 ± 3.76	12.01 ± 7.46	10.01 ± 4.98	1.2108 ± 0.3042
Trexplorer Super	**55.27 ± 3.00**	**33.99 ± 3.34**	**39.46 ± 1.93**	**0.5627 ± 0.0141**

Table 3. Comparison of different models using branch-level and graph-level metrics on the Synthetic, ATM'22, and Parse 2022 datasets.

Model	Branch-level			Graph-level (MAE)	
	Precision(%)↑	Recall(%)↑	F1(%)↑	Betti-0↓	Betti-1↓
Synthetic Dataset					
Vesselformer	12.51 ± 0.64	27.08 ± 2.85	15.95 ± 0.36	81.7 ± 16.8	653.5 ± 138.7
Trexplorer	19.20 ± 7.31	64.91 ± 3.68	26.26 ± 7.18	**0.000 ± 0.0**	**0.000 ± 0.0**
Trexplorer Super	**96.01 ± 4.42**	**67.52 ± 2.94**	**77.12 ± 1.59**	**0.000 ± 0.0**	**0.000 ± 0.0**
ATM'22 Dataset					
Vesselformer	1.35 ± 0.19	3.67 ± 0.41	1.95 ± 0.25	312.5 ± 25.1	180.4 ± 35.9
Trexplorer	0.03 ± 0.01	0.14 ± 0.04	0.05 ± 0.02	**0.00 ± 0.0**	**0.00 ± 0.0**
Trexplorer Super	**45.15 ± 2.39**	**42.23 ± 1.54**	**41.15 ± 1.28**	**0.00 ± 0.0**	**0.00 ± 0.0**
Parse 2022 Dataset					
Vesselformer	2.32 ± 0.20	1.89 ± 0.18	1.99 ± 0.16	410.1 ± 23.9	246.7 ± 78.1
Trexplorer	3.40 ± 1.41	5.36 ± 3.50	3.71 ± 1.91	**0.00 ± 0.0**	**0.00 ± 0.0**
Trexplorer Super	**35.45 ± 2.89**	**20.09 ± 1.86**	**23.46 ± 1.09**	**0.00 ± 0.0**	**0.00 ± 0.0**

whereas Vesselformer struggles, predicting multiple disconnected components and cycles.

The performance of Trexplorer Super is impacted by certain cases where it fails to track the centerline, likely due to unseen intensity variations in the input volumes. Applying image augmentation could help address this issue and further improve performance.

3.5 Ablations

We conduct an ablation study on the ATM'22 dataset to evaluate the impact of our key modifications to Trexplorer: Super Trajectory Training, Focal Cross

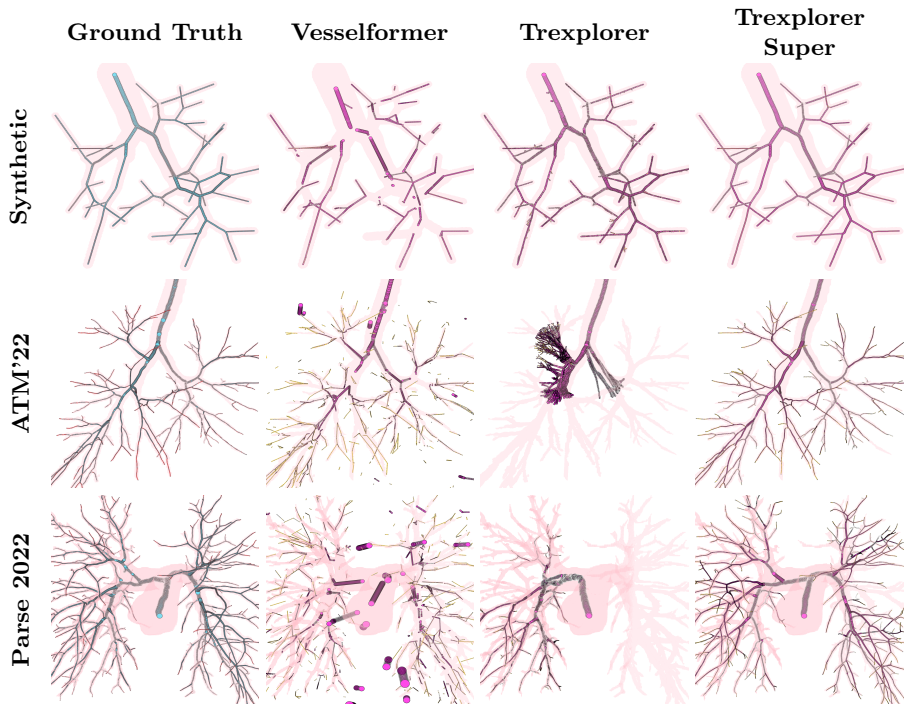

Fig. 2. Visual comparison between the ground truth and the predictions from Vesselformer, Trexplorer, and Trexplorer Super on one sample from each dataset. Centerline marker size is proportional to the radius at each node.

Attention, and Target Augmentation. Table 4 reports the mean and standard deviation of point-level metrics, including precision, recall, and F1-score, averaged over three runs for each ablation. The results highlight Super Trajectory Training as the most critical improvement, allowing Trexplorer to perform effectively on real data. Focal Cross Attention further enhances performance by enabling the feature extractor to condense relevant information in the focal region. Target Augmentation improves both recall and precision by reducing duplicate predictions and enhancing the detection of new branches, leading to more complete branch reconstructions.

4 Conclusion

We present major improvements to Trexplorer, enhancing its accuracy, robustness, and completeness. Additionally, we introduce three new datasets and conduct a comprehensive evaluation, demonstrating that our model outperforms two state-of-the-art baselines. While the results are promising, challenges remain, particularly on the Parse 2022 dataset. Future work includes leveraging a

Table 4. Evaluation of the novel key components, Super Trajectory Training (STT), Focal Cross Attention (FCA), and Target Augmentation (TA) incorporated in the Trexplorer framework by Trexplorer Super for the ATM'22 Dataset.

Index	Novel Key Components			Point-level		
	STT	FCA	TA	Precision(%)↑	Recall(%)↑	F1(%)↑
1				3.07 ± 0.08	3.59 ± 0.21	3.17 ± 0.11
2		✓		36.57 ± 54.94	3.66 ± 3.20	2.56 ± 2.19
3	✓			44.87 ± 10.20	27.43 ± 7.69	33.01 ± 8.70
4	✓	✓		61.76 ± 2.23	54.72 ± 7.45	53.91 ± 5.05
5	✓	✓	✓	**67.82 ± 1.04**	**61.20 ± 2.50**	**60.66 ± 0.64**

larger pretrained feature extractor and integrating advanced DETR variants or recurrent architectures, such as LSTMs, for further refinement.

Acknowledgments. Compute and storage resources were provided by NAISS (grant 2022-06725, Swedish Research Council) and the Berzelius system at the National Supercomputer Centre, funded by the Knut and Alice Wallenberg Foundation.

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

References

1. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
2. Choi, A.D., Marques, H., Kumar, V., Griffin, W.F., Rahban, H., Karlsberg, R.P., Zeman, R.K., Katz, R.J., Earls, J.P.: Ct evaluation by artificial intelligence for atherosclerosis, stenosis and vascular morphology(clarify): A multi-center, international study. Journal of Cardiovascular Computed Tomography **15**(6), 470–476 (2021)
3. Community, S.D.: 3d slicer (2025), `https://www.slicer.org/`, accessed: 2025-02-24
4. F., S.A., et al.: Vmtk: A toolkit for computational vascular modeling. Journal of Computational Science **23**, 59–69 (2018). `https://doi.org/10.1016/j.jocs.2018.02.003`, `https://vmtk.org/`, accessed: 2025-02-24
5. Huang, D., Tang, W., Ding, Y., Wan, T., Chen, Y.: An interactive 3d preoperative planning and training system for minimally invasive vascular surgery. In: 2011 12th International Conference on Computer-Aided Design and Computer Graphics. pp. 443–449. IEEE (2011)
6. Khan, Z., Ngo, J.P., Le, B., Evans, R.G., Pearson, J.T., Gardiner, B.S., Smith, D.W.: Three-dimensional morphometric analysis of the renal vasculature. American Journal of Physiology-Renal Physiology **314**(5), F715–F725 (2018)
7. Li, Z., Xia, Q., Hu, Z., Wang, W., Xu, L., Zhang, S.: A deep reinforced tree-traversal agent for coronary artery centerline extraction. In: Medical Image Computing

and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part V 24. pp. 418–428. Springer (2021)
8. Luo, G., Wang, K., Liu, J., Li, S., Liang, X., Li, X., Gan, S., Wang, W., Dong, S., Wang, W., et al.: Efficient automatic segmentation for multi-level pulmonary arteries: The parse challenge. arXiv preprint arXiv:2304.03708 (2023)
9. Miraucourt, O., Salmon, S., Szopos, M., Thiriet, M.: Blood flow in the cerebral venous system: modeling and simulation. Computer methods in biomechanics and biomedical engineering **20**(5), 471–482 (2017)
10. Naeem, R., Hagerman, D., Svensson, L., Kahl, F.: Trexplorer: Recurrent detr for topologically correct tree centerline tracking. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 744–754. Springer (2024)
11. Prabhakar, C., Shit, S., Paetzold, J.C., Ezhov, I., Koner, R., Li, H., Kofler, F.S., Menze, B.: Vesselformer: Towards complete 3d vessel graph generation from images. In: Medical Imaging with Deep Learning. pp. 320–331. PMLR (2024)
12. Qin, Y., Chen, M., Zheng, H., Gu, Y., Shen, M., Yang, J., Huang, X., Zhu, Y.M., Yang, G.Z.: Airwaynet: a voxel-connectivity aware approach for accurate airway segmentation using convolutional neural networks. In: International conference on medical image computing and computer-assisted intervention. pp. 212–220. Springer (2019)
13. Roh, B., Shin, J., Shin, W., Kim, S.: Sparse detr: Efficient end-to-end object detection with learnable sparsity. arXiv preprint arXiv:2111.14330 (2021)
14. Sexton, Z.A.: Synthetic vascular toolkit (2023), `https://github.com/zasexton/Synthetic-Vascular-Toolkit`, accessed: 2025-02-24
15. Sexton, Z.A., Hudson, A.R., Herrmann, J.E., Shiwarski, D.J., Pham, J., Szafron, J.M., Wu, S.M., Skylar-Scott, M., Feinberg, A.W., Marsden, A.: Rapid model-guided design of organ-scale synthetic vasculature for biomanufacturing. ArXiv pp. arXiv–2308 (2023)
16. Shit, S., Koner, R., Wittmann, B., Paetzold, J., Ezhov, I., Li, H., Pan, J., Sharifzadeh, S., Kaissis, G., Tresp, V., et al.: Relationformer: A unified framework for image-to-graph generation. In: European Conference on Computer Vision. pp. 422–439. Springer (2022)
17. Silversmith, W., Bae, J.A., Li, P.H., Wilson, A.: Kimimaro: Skeletonize densely labeled 3d image segmentations (2021). `https://doi.org/10.5281/zenodo.5539913`, `https://doi.org/10.5281/zenodo.5539913`, accessed: 2021-09-29
18. Tetteh, G., Efremov, V., Forkert, N.D., Schneider, M., Kirschke, J., Weber, B., Zimmer, C., Piraud, M., Menze, B.H.: Deepvesselnet: Vessel segmentation, centerline prediction, and bifurcation detection in 3-d angiographic volumes. Frontiers in Neuroscience **14**, 1285 (2020)
19. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
20. Yu, W., Zheng, H., Zhang, M., Zhang, H., Sun, J., Yang, J.: Break: Bronchi reconstruction by geodesic transformation and skeleton embedding. In: 2022 IEEE 19th international symposium on biomedical imaging (ISBI). pp. 1–5. IEEE (2022)
21. Zhang, M., Wu, Y., Zhang, H., Qin, Y., Zheng, H., Tang, W., Arnold, C., Pei, C., Yu, P., Nan, Y., et al.: Multi-site, multi-domain airway tree modeling. Medical Image Analysis **90**, 102957 (2023)

22. Zhang, M., Zhang, H., Yang, G.Z., Gu, Y.: Cfda: collaborative feature disentanglement and augmentation for pulmonary airway tree modeling of covid-19 cts. In: International conference on medical image computing and computer-assisted intervention. pp. 506–516. Springer (2022)
23. Zhang, Y., Luo, G., Wang, W., Wang, K.: Branch-aware double dqn for centerline extraction in coronary ct angiography. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23. pp. 35–44. Springer (2020)
24. Zheng, D., Dong, W., Hu, H., Chen, X., Wang, Y.: Less is more: Focus attention for efficient detr (2023)
25. Zheng, H., Qin, Y., Gu, Y., Xie, F., Yang, J., Sun, J., Yang, G.Z.: Alleviating class-wise gradient imbalance for pulmonary airway segmentation. IEEE transactions on medical imaging **40**(9), 2452–2462 (2021)
26. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)