

Anatomy-Aware Prediction of Bronchoscopic Accessibility from 3D CT

Linkai Peng¹, Cuiling Sun², Bin Wang¹, Jamie Rowell³, Catherine Gao³, Oyku Ikizgul⁴, Eminenur Sentasci⁵, Andrea Bejar⁵, Halil Ertugrul Aktas⁵, Gorkem Durak⁵, Momen Wahidi³, Christopher Kapp³, and Ulas Bagci^{5†}

¹ Department of Electrical and Computer Engineering, Northwestern University

² Department of Computer Science, Northwestern University

³ Department of Medicine, Northwestern University

⁴ Department of Radiology, Istanbul Faculty of Medicine

⁵ Department of Radiology, Northwestern University

† Corresponding author: ulas.bagci@northwestern.edu

Abstract. Pre-operative planning for bronchoscopy is critical for the diagnosis of lung lesions. Current accessibility assessment relies on subjective manual inspection of CT scans, which is time-consuming and prone to inter-observer variability. In this paper, we formalize bronchoscopy accessibility prediction as a novel supervised learning task and present the first end-to-end framework to address it. We propose an **Anatomy-Aware Mixture-of-Experts (MoE)** model that integrates specialized modules: a CT Expert for local morphological features, a Lobe Expert for anatomical priors, and a Path Geometry Expert that encodes the sequential constraints of the bronchial tree. To support this task, we curated **the first clinical dataset of 438 cases** with pre-operative CT scans and documented procedural outcomes. Experimental results demonstrate that our method achieves an AUROC of 0.8052, significantly outperforming both state-of-the-art baselines and experienced human experts. This work establishes a new benchmark for computer-aided interventional planning in pulmonary medicine. Our data and code will be publicly available at <https://nubagcilab.github.io/BronchoAccess/>.

Keywords: Mixture of Expert · Bronchoscopy Accessibility · Clinical Planning.

1 Introduction

Bronchoscopy is widely used for the diagnosis of pulmonary lesions. Although modern navigation systems assist intra-procedural guidance [6,12], pre-operative accessibility assessment remains largely subjective. As demonstrated in our benchmark, expert bronchoscopists predict a diagnostic procedure with an AUROC of only 0.5661. This performance highlights the intrinsic difficulty of the task.

Despite its clinical relevance, bronchoscopy accessibility prediction remains underexplored. Prior work has focused on airway segmentation [23,25,26], lesion detection [19], or path planning [5]. While these tasks provide an important

basis, they do not directly address procedural accessibility. Existing attempts at outcome prediction rely on handcrafted geometric descriptors and statistical association models without end-to-end supervision [18], failing to capture non-linear interactions between local morphology, anatomical location, and airway complexity. A standardized benchmark for this task is also absent.

To bridge this gap, we first formalize bronchoscopy accessibility prediction as a supervised learning task and curate the first dataset dedicated to this problem. We retrospectively collected 438 clinically verified cases from three medical centers, each with documented procedural outcomes following bronchoscopy for lung lesions. This multi-center cohort provides diverse anatomical variability and realistic clinical distributions.

Building upon this task definition, we propose an anatomy-aware mixture-of-experts (MoE) [20] framework that decomposes accessibility prediction into three components. The CT Expert captures local volumetric context around the lesion. The Lobe Expert encodes anatomical priors reflecting regional variations. The Path Geometry Expert models cumulative geometric constraints along the airway trajectory. These representations are adaptively fused through a gating mechanism that performs case-dependent expert weighting.

To enable systematic evaluation, we further establish the first comprehensive benchmark for bronchoscopy accessibility prediction. The benchmark includes traditional machine learning methods, CT-based deep networks, graph neural networks, and human expert assessment. This benchmark provides a unified evaluation framework for future research on data-driven procedural planning.

Our contributions are three-fold:

1. To address this severely under-studied clinical challenge, we formalize bronchoscopy accessibility prediction as an end-to-end supervised deep learning task. To the best of our knowledge, this work presents the first deep learning application dedicated to predicting functional procedural outcomes directly from pre-operative CT scans.
2. We introduce an anatomy-aware MoE architecture that integrates volumetric imaging, anatomical priors, and airway geometry within a unified model.
3. We establish the first comprehensive benchmark for bronchoscopy accessibility prediction using 438 clinically verified cases and evaluate our framework against diverse computational baselines and human expert assessments.

2 Method

Given CT volume $X \in \mathbb{R}^{H \times W \times D}$, airway tree graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, target lesion centroid $\mathbf{c}_N \in \mathbb{R}^3$ and lobe identifier $L \in \{1, \dots, 5\}$, we learn $\mathcal{F} : \{X, \mathcal{G}, \mathbf{c}_N, L\} \rightarrow y$, where $y \in \{0, 1\}$ denotes bronchoscopy accessibility. Our framework decomposes this into three components as shown in Fig. 1.

2.1 Expert Modules

CT Expert The local structural relationship between the target lesion and the surrounding airway tree plays a central role in bronchoscopy accessibility. CT




Fig. 1. Overview of the proposed anatomy-aware mixture-of-experts framework for bronchoscopy accessibility prediction.

Expert is formulated to capture the local morphological and spatial relationship between the target lesion and its surrounding airway structures.

We extract a multi-channel 3D patch x comprising the raw CT intensities, a binary lesion mask, and a 3D Euclidean Distance Transform (EDT) [17] of the airway. The lesions are first automatically detected and subsequently verified to ensure consistency with the procedural target. The EDT provides continuous spatial gradients that stabilize learning compared to sparse airway masks. In this way, the model receives an explicit geometric prior reflecting airway accessibility rather than relying solely on implicit learning from intensity patterns.

The multi-channel patch is processed by a custom 3D convolutional encoder. It consists of 4 consecutive residual blocks [8] integrated with squeeze-and-excitation (SE) [10] channel attention. After convolution, a global average pooling layer aggregates the spatial features and a linear layer projects it into a latent embedding $z_{ct} \in \mathbb{R}^d$, where d is the joint feature space dimension.

Lobe Expert Bronchoscopy difficulty varies across pulmonary lobes due to anatomical differences in branching depth and orientation. To incorporate this anatomical prior, we encode the lobe identity L as a categorical variable and map it into a learnable embedding space via an embedding layer followed by a two-layer MLP: $z_{lobe} = \text{MLP}(\Psi(L))$, $z_{lobe} \in \mathbb{R}^d$. This embedding serves as a global anatomical prior that regularizes the final decision. By encoding lobe identity as a continuous representation, the model can learn structured similarities between lobes rather than treating them as independent categories.

Path Geometry Expert Accessibility is constrained by cumulative geometric resistance along the airway trajectory from the trachea to the target region. These constraints are sequential and longitudinal in nature, which motivates explicit modeling of the airway trajectory. While the airway tree is naturally

a graph, once the navigation path is selected, the data along it forms a one-dimensional ordered sequence of geometric descriptors; we therefore adopt a sequence model rather than a graph neural network. We represent the airway tree as a graph $\mathcal{G}$ derived from the segmented airway and extract the candidate navigation path toward the lesion. For cases with multiple plausible routes, the path with the minimal Euclidean distance to the lesion centroid is selected. Nodes along this path are ordered according to anatomical generation to form a proximal-to-distal sequence. To ensure robustness to segmentation noise, we retain only the largest connected component and remove edges shorter than a predefined physical threshold.

At each node, we compute a set of geometric and topological descriptors, including airway diameter, branching index, curvature, torsion, and relative path length. The ordered feature sequence forms a representation $S = \{\mathbf{s}_1, \dots, \mathbf{s}_T\}$, where T differs across patients depending on airway depth.

To model longitudinal constraints, we apply a 1-layer bidirectional Gated Recurrent Unit (GRU) [4] over the sequence. The GRU is configured with a hidden size of 128 per direction. The final hidden states from both directions are concatenated and then projected to the joint feature space to obtain a path representation z_{path} . By explicitly modeling the airway as a sequence, this expert captures progressive narrowing, branching accumulation, and distal complexity. This design provides the framework with a topological understanding of accessibility that is complementary to the visual patterns captured by the CT Expert.

2.2 Anatomy-Aware Adaptive Expert Fusion

To handle cross-case heterogeneity, we employ an adaptive mixture-of-experts fusion. A gating network assigns importance weights $\mathbf{w} = [w_{ct}, w_{lobe}, w_{path}]^\top$ to calculate the final fused representation $Z = \sum w_i z_i$. The gating network conditions on case-level anatomical context: target lesion position p , lobe embedding $\Psi(L)$, and graph-level airway statistics g , producing weights via $w = \text{Softmax}(\text{MLP}([p; \Psi(L); g]))$. It is implemented as a three-layer MLP with ReLU activations and a final Softmax layer. The fused representation Z is then passed through a classification head to produce the probability of bronchoscopy accessibility $\hat{y} \in [0, 1]$. The entire system is trained using the binary cross-entropy (BCE) loss $\mathcal{L}_{bce}$. To encourage balanced expert utilization, we add an entropy regularization term on the gating weights:

$$\mathcal{L}_{ent} = -\frac{1}{M} \sum_{i=1}^M \sum_{j \in \{ct, lobe, path\}} w_{i,j} \log(w_{i,j} + \epsilon), \quad (1)$$

where M denotes the number of samples and ϵ is a small constant for numerical stability. We encourage higher entropy in the gating distribution to avoid degenerate solutions where a single expert dominates across all cases. The total training loss $\mathcal{L}_{total}$ is formulated as: $\mathcal{L}_{total} = \mathcal{L}_{bce} - \lambda \mathcal{L}_{ent}$, where λ is a hyperparameter controlling the regularization strength.

3 Experiments

3.1 Dataset

To fully evaluate the proposed framework, we curated a retrospective multi-center cohort consisting of 438 patients who underwent bronchoscopy for lung lesions. The dataset comprises cases from three clinical centers, including 279 successful procedures and 159 failed attempts. A procedure is labeled successful when diagnostic tissue is confirmed by histopathological analysis of the biopsy sample, which is an objective outcome based on diagnostic yield. Accessibility labels are derived from documented clinical outcomes. All procedures were performed using the same platform to ensure consistency in primary navigation technology across cases. To our knowledge, this is currently the largest dataset reported for outcome-driven bronchoscopy accessibility prediction. The dataset also involves multiple operators and varying device configurations. Although these procedural factors were not explicitly modeled, the multi-center design introduces heterogeneity that reduces systematic bias from a single operator or device type. To ensure robust evaluation, we employ 5-fold cross-validation with patient-level splits.

3.2 Implementations and Compared Methods

For the CT Expert, 3D patches of $128 \times 128 \times 32$ are extracted around the target lesion. Intensity values were clipped to the lung window range and normalized. We used RetinaNet [1,3,14] for lung lesion detection and Naviairway [23] for airway segmentation. All models are trained using AdamW with a learning rate of 1×10^{-3} , batch size 16, and 100 epochs on an NVIDIA A100 GPU. To provide clinical context, 22 clinicians independently assessed accessibility from CT volumes and target locations. Three-level difficulty ratings were mapped to probabilistic scores (0.9, 0.7, 0.4) to compute AUROC and PRAUC. The coefficient λ is set to 0.01. Dropout and weight decay were used during training to prevent overfitting. We compare against conventional machine learning models trained on handcrafted geometric descriptors, CT-based deep networks (ResNet-18 [8], DenseNet-121 [11], EfficientNet-b0 [21], ConvNeXt [16], and Swin Transformer [15]), and graph neural networks (GCN [13], GAT [22], GIN [24], and GraphSAGE [7]) operating on airway topology.

4 Results

4.1 Quantitative Results

Comparison with Diverse Baselines Table 1 presents the overall comparison across various methods. Conventional machine learning models trained solely on handcrafted geometric features show limited discriminative ability. CT-based deep networks achieve moderate performance, with the best AUROC reaching 0.6333. Graph neural networks slightly improve structural modeling, achieving

up to 0.6483 AUROC. Human experts achieve an AUROC of 0.5661, indicating that visual CT inspection alone is insufficient for reliable accessibility estimation.

The proposed anatomy-aware mixture-of-experts model achieves an AUROC of 0.8052 and PR-AUC of 0.8496, significantly outperforming computational baselines and human assessment. These results demonstrate the benefit of explicitly integrating airway trajectory modeling with local imaging cues. Importantly, the proposed model maintains competitive computational efficiency. With 3.83M parameters, it is more lightweight than most CT-based backbone networks while achieving superior predictive performance. This efficiency supports potential integration into pre-procedural planning workflows.

Table 1. Performance comparison of the proposed method against various baseline models. Human assessments were provided by 22 clinicians based on CT volumes and target lesion locations. [†] indicates statistically significant difference compared with the proposed method (DeLong test, $p < 0.05$).

Input	Method	Acc.	F1	AUROC	PRAUC	Params
Geometric Features	Logistic Regression[9]	0.6136	0.7536	0.5128	0.6897	N/A
	Random Forest[2]	0.5682	0.7031	0.4590	0.5785	N/A
CT Volumes	Densenet121[11] [†]	0.5909	0.7313	0.5301	0.6543	11.24M
	SwinTransformer[15] [†]	0.6136	0.7499	0.5418	0.8069	38.50M
	Efficientnetb0[21] [†]	0.6022	0.7517	0.5725	0.7459	4.69M
	ConvNeXt[16] [†]	0.6023	0.7482	0.5982	0.7750	31.30M
	Resnet18[8] [†]	0.5795	0.6185	0.6333	0.7700	33.16M
Airway Graph	GAT[22] [†]	0.6386	0.7794	0.5767	0.7115	52.35K
	GIN[24] [†]	0.6265	0.7257	0.5774	0.6782	101.12K
	GraphSage[7] [†]	0.6203	0.7541	0.5897	0.7129	86.14K
	GCN[13] [†]	0.6329	0.7752	0.6483	0.7666	51.59K
Human Experts	N/A	0.5272	0.5423	0.5661	0.7434	N/A
Hybrid	Ours	0.8068	0.8595	0.8052	0.8496	3.83 M

Comparison Under Identical CT Inputs To ensure that the performance gain does not arise from richer input representations alone, Table 2 compares our model with CT-based deep networks trained on identical multi-channel inputs, including CT intensity, airway distance maps, and lesion masks. Even under this controlled setting, our model maintains a clear advantage. The strongest CNN baseline achieves 0.7310 AUROC, while the proposed method reaches 0.8052 AUROC. This result indicates that architectural decomposition and adaptive expert fusion account for the performance improvement rather than input augmentation alone.

Ablation Study Table 3 evaluates the contribution of each expert and their combinations. Among single branches, the CT Expert achieves the strongest performance (AUROC 0.7382), confirming that local volumetric context provides a

solid baseline. The Path Geometry Expert (0.6040 AUROC) captures meaningful structural constraints along the airway trajectory but remains insufficient alone. The Lobe Expert performs near chance level (0.5513 AUROC), consistent with its role as a global anatomical prior rather than a standalone predictor. Pairwise

Table 2. Comparison with CT-based deep networks trained using multi-channel volumetric inputs (CT intensity, airway distance transform, and lesion mask). [†] indicates statistically significant difference compared with the proposed method (DeLong test, $p < 0.05$).

Method	Acc.	F1	AUROC	PRAUC
SwinTransformer [†] [15]	0.6250	0.7626	0.5243	0.6698
ConvNeXt [†] [16]	0.6477	0.7438	0.6629	0.7057
Densenet121 [†] [11]	0.6590	0.7058	0.6640	0.7611
Resnet18 [†] [8]	0.6818	0.7586	0.7109	0.7834
Efficientnetb0 [21]	0.6704	0.7819	0.7310	0.8269
Ours	0.8068	0.8595	0.8052	0.8496

combinations consistently improve over individual modules. Combining CT and Path Experts increases AUROC to 0.7756, demonstrating the complementary effect of trajectory modeling. The full model integrating all three experts achieves the best performance (0.8052 AUROC). We also observe that CT Expert receives higher weights for centrally located lesions, while Path Geometry Expert weights increase for distal and lower-lobe lesions, suggesting anatomically consistent specialization. The gating weights, $\mathbf{w} = [w_{ct}, w_{lobe}, w_{path}]^\top$, thus provide a built-in case-level attribution: for any prediction, $\mathbf{w}$ indicates which anatomical factor drives the decision, offering interpretability without requiring post-hoc explanation methods. These results indicate that bronchoscopy accessibility depends on the joint modeling of local morphology, anatomical context, and cumulative airway geometry rather than any single modality alone.

Table 3. Ablation study evaluating the contribution of each expert module. The full model integrates all experts through adaptive gating.

Lobe Expert	Path Expert	CT Expert	Acc.	F1	AUROC	PRAUC
✓	-	-	0.6190	0.7647	0.5513	0.7030
-	✓	-	0.6310	0.7704	0.6040	0.6665
-	-	✓	0.7500	0.8035	0.7382	0.8028
✓	✓	-	0.7045	0.7968	0.5960	0.6857
✓	-	✓	0.7841	0.8504	0.7477	0.8141
-	✓	✓	0.7500	0.8358	0.7756	0.8030
✓	✓	✓	0.8068	0.8595	0.8052	0.8496

4.2 Qualitative Analysis

To further illustrate model behavior, Fig. 2 presents representative cases of true positive, false positive, false negative, and true negative predictions. The red marker indicates the target lesion location, and the airway tree is rendered to visualize the navigational trajectory. The true positive case shows a lesion near




Fig. 2. Representative qualitative examples illustrating true positive (TP), false positive (FP), false negative (FN), and true negative (TN) predictions. Airway trees are visualized together with target lesion locations (red markers).

a well-connected airway branch with sufficient diameter. The false positive case shows a lesion anatomically close to the distal bronchi but with abrupt branching changes along the path, where the CT Expert may overweight local proximity in such cases. The false negative involves a peripheral lesion with limited distal branches, where successful sampling may still occur in practice due to operator expertise or ultra-thin bronchoscopes. The true negative shows a deeply peripheral lesion with sparse connectivity. Overall, accessibility depends on both local lesion–airway adjacency and cumulative geometric constraints; clinical execution factors beyond static anatomy also influence outcomes.

5 Conclusion

To the best of our knowledge, this is the first study to systematically formalize bronchoscopy accessibility prediction as a supervised learning problem and to benchmark it against both computational baselines and human expert assessment. The results reveal that this clinically relevant task remains under-explored and difficult under conventional visual evaluation paradigms.

We presented the first learning-based framework for predicting bronchoscopy accessibility directly from pre-procedural CT. The proposed anatomy-aware mixture-of-experts model decomposes accessibility prediction into complementary com-

ponents. It integrates local volumetric context, global anatomical priors, and sequential airway path geometry within a unified end-to-end architecture.

The experimental results demonstrate that our approach significantly outperforms both state-of-the-art single-modality baselines and experienced human experts. Notably, the model achieves these gains with a lightweight architecture, maintaining low parameter count and computational cost.

By bridging airway topology and functional procedural outcome prediction, this work moves beyond connectivity analysis toward quantitative accessibility assessment. We anticipate that such anatomy-aware modeling may support pre-procedural planning and improve patient selection in bronchoscopy-guided diagnosis.

Acknowledgments. This work was partially supported by NIH R01-HL171376.

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

References

1. Baumgartner, M., Jäger, P.F., Isensee, F., Maier-Hein, K.H.: nndetection: a self-configuring method for medical object detection. In: International conference on medical image computing and computer-assisted intervention. pp. 530–539. Springer (2021)
2. Breiman, L.: Random forests. *Machine learning* **45**(1), 5–32 (2001)
3. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al.: Monai: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701* (2022)
4. Cho, K., Van Merriënboer, B., Gulcehre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder–decoder for statistical machine translation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). pp. 1724–1734 (2014)
5. Ciobirca, C., Lango, T., Gruionu, G., Leira, H.O., Gruionu, L., Pastrama, S.: A new procedure for automatic path planning in bronchoscopy. *Materials Today: Proceedings* **5**(13), 26513–26518 (2018)
6. Folch, E.E., Pritchett, M.A., Nead, M.A., Bowling, M.R., Murgu, S.D., Krinsky, W.S., Murillo, B.A., LeMense, G.P., Minnich, D.J., Bansal, S., et al.: Electromagnetic navigation bronchoscopy for peripheral pulmonary lesions: one-year results of the prospective, multicenter navigate study. *Journal of Thoracic Oncology* **14**(3), 445–458 (2019)
7. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. *Advances in neural information processing systems* **30** (2017)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
9. Hosmer Jr, D.W., Lemeshow, S., Sturdivant, R.X.: Applied logistic regression. John Wiley & Sons (2013)
10. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7132–7141 (2018)

11. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)
12. Khandhar, S.J., Bowling, M.R., Flandes, J., Gildea, T.R., Hood, K.L., Krinsky, W.S., Minnich, D.J., Murgu, S.D., Pritchett, M., Toloza, E.M., et al.: Electromagnetic navigation bronchoscopy to access lung lesions in 1,000 subjects: first results of the prospective, multicenter navigate study. *BMC Pulmonary Medicine* **17**(1), 59 (2017)
13. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
14. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
15. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
16. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11976–11986 (2022)
17. Meijster, A., Roerdink, J.B., Hesselink, W.H.: A general algorithm for computing distance transforms in linear time. In: Mathematical Morphology and its applications to image and signal processing, pp. 331–340. Springer (2000)
18. Naito, M., Masaki, F., Lisk, R., Tsukada, H., Hata, N.: Predicting reachability to peripheral lesions in transbronchial biopsies using ct-derived geometrical attributes of the bronchial route. *International Journal of Computer Assisted Radiology and Surgery* **18**(2), 247–255 (2023)
19. Nasrullah, N., Sang, J., Alam, M.S., Mateen, M., Cai, B., Hu, H.: Automated lung nodule detection and classification using deep learning combined with multiple strategies. *Sensors* **19**(17), 3722 (2019)
20. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., Dean, J.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. arXiv preprint arXiv:1701.06538 (2017)
21. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International conference on machine learning. pp. 6105–6114. PMLR (2019)
22. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017)
23. Wang, A., Tam, T.C.C., Poon, H.M., Yu, K.C., Lee, W.N.: Naviairway: a bronchiole-sensitive deep learning-based airway segmentation pipeline. arXiv preprint arXiv:2203.04294 (2022)
24. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? arXiv preprint arXiv:1810.00826 (2018)
25. Yang, X., Chen, L., Zheng, Y., Ma, L., Chen, F., Ning, G., Liao, H.: Airway segmentation based on topological structure enhancement using multi-task learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 86–95. Springer (2024)
26. 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)