

Semantically Consistent Discrete Diffusion for 3D Biological Graph Modeling

Chinmay Prabhakar^{*1,2} , Suprosanna Shit^{*1,2} , Tamaz Amiranashvili^{1,3} ,
Hongwei Bran Li⁴ , and Bjoern Menze¹ 

¹ Department of Quantitative Biomedicine, University of Zurich, Switzerland

² ETH AI Center, ETH Zurich, Switzerland

³ Department of Computer Science, Technical University of Munich, Germany

⁴ Athinoula A. Martinos Center, Harvard Medical School, USA

chinmay.prabhakar@uzh.ch

Abstract. 3D spatial graphs play a crucial role in biological and clinical research by modeling anatomical networks such as blood vessels, neurons, and airways. However, generating 3D biological graphs while maintaining anatomical validity remains challenging, a key limitation of existing diffusion-based methods. In this work, we propose a novel 3D biological graph generation method that adheres to structural and semantic plausibility conditions. We achieve this by using a novel projection operator during sampling that stochastically fixes inconsistencies. Further, we adopt a superior edge-deletion-based noising procedure suitable for sparse biological graphs. Our method demonstrates superior performance on two real-world datasets, human circle of Willis and lung airways, compared to previous approaches. Importantly, we demonstrate that the generated samples significantly enhance downstream graph labeling performance. Furthermore, we show that our generative model is a reasonable out-of-the-box link predictor.

Keywords: Discrete Diffusion · Vessel Graph · Airways.

1 Introduction

3D spatial graphs are an important representation for the study of biological graphs, including vascular networks [12], neural connectivity [25], airway trees [11], lymphatic networks [18], and bone microarchitecture [13]. These graphs are obtained from the segmentation [23] of imaging modalities. Several works [15, 10] have studied the characteristics of each of these graphs and their relevance as biomarkers for detecting abnormalities. Learning the distribution of biological graphs is crucial not only because it can serve as a way to generate synthetic samples but also because it can be used as a strong prior for downstream tasks such as link prediction, labeling, etc. However, learning the distribution of biological graphs is challenging due to the difficulty of enforcing anatomical

* Contributed equally

plausibility, both in terms of structural constraints (e.g., avoiding cycles in airway trees) and semantic coherence (e.g., preventing anatomically inconsistent edge labels between neighboring nodes).

Biological graph modeling has been studied using stochastic methods [16] or topology optimization [6]. These models do not explicitly learn the underlying distribution but are rather sampled from a heuristic model. Recently, generative models [14] have been used to model vascular graphs. However, there is no guarantee of generating anatomically plausible graphs, often leading to reliance on sub-optimal post-processing or inefficient sample rejection. Further, the discrete noising model considered in [14] changes the ground truth edge class arbitrarily to a different edge class. This noising procedure is inefficient since the diffusion model wastes resources on discarding a significant amount of implausible long-range edges. Recently, Madeira et al. [9] introduced an alternative edge noising model that randomly deletes existing edges, which mitigates the problem of the appearance of implausible edges in the noising process. Furthermore, they provided theoretical results showing that if a desired graph property remains invariant under a noising procedure, it can be enforced by a hard projection operator during sampling. This allows one to generate graphs that maintain structural properties, such as a tree-like topology. However, biological graphs often adhere to strict *semantic* constraints on top of structural ones. For instance, a basal artery (BA) must not be directly connected to the anterior cerebral artery (ACA). This represents an important biological prior, which cannot be imposed by projections in [9]. In this work, we go beyond structural constraints and introduce a projector module that stochastically corrects the intermediate graph by removing edges that violate structural constraints (e.g., causing cycles in the airway tree) *and/or* semantic constraints (e.g., BA and ACA are direct neighbors). In summary, our contributions are threefold:

1. We propose a novel label-consistent projector that, during sampling, prevents conflicting neighboring edge labels that are anatomically invalid, thus ensuring anatomically plausible graphs- as demonstrated in our experiments in two real-world datasets, human circle of Willis and lung airway.
2. We adopt an edge deletion noising model for the discrete diffusion, which is suitable for sparse biological graphs.
3. We validate the efficacy of generated samples from our proposed generation to train a vessel labeler that significantly boosts its accuracy. Furthermore, we demonstrate excellent out-of-the-box emerging link prediction properties of our graph generator in the airway dataset. Code is available at ⁵

2 Related Literature

Biological graph generation is a complex modeling task due to anatomical variations, making specialized heuristics-driven techniques very popular. For instance, stochastic generation [17,16] and topology optimization [6] rely on heuristics to generate realistic vessel trees, while Nousias et al. [11] applied heuristics

⁵ https://github.com/chinmay5/semantically_consistent_graph_generation




Fig. 1. Our method first generates node coordinates using a point cloud denoiser. Subsequently, the edge denoiser is trained with an edge-deletion noise model. During inference, the edge denoiser suggests a set of new edges. Our novel semantic consistent projector discards an edge if none of k resampled edge labels meet the validity criteria.

to generate airway trees. Recently, data-driven methods have been proposed [7, 14, 25, 11] to avoid such heuristics. Among these, Prabhakar and Shit et al. [14] propose a model capable of generating complex vascular graphs. Their method is extendable to other spatial structures. However, it can generate semantically invalid graphs, leading to data inconsistencies.

Discrete diffusion models excel in modeling discrete data and are widely used in molecule generation [20, 21], protein structure design [4], and language modeling [1, 8]. However, they often fail to enforce domain-specific constraints, leading to invalid samples. To address this, Sharma et al. [19] introduce plug-and-play closed-form guidance for properties such as edge count, degree, and triangle count, etc., but it does not guarantee constraint satisfaction. Recently, Madeira et al. [9] proposes hard-projection criteria to enforce validity. Note that their method does not consider edge labels and focuses only on structural constraints (tree structure, node-degree, etc).

3 Methodology

Problem Statement: Our goal is to learn the distribution $p(G)$ of 3D spatial graphs $G := (\mathbf{X}, \mathbf{E})$, where $\mathbf{X} \in \mathbb{R}^{n \times 3}$ represents coordinates of n nodes, $\mathbf{E} \in \mathbb{R}^{n \times n \times c}$ is adjacency matrix containing one-hot encoding of categorical edge labels indicating vessel class and E is the corresponding set of edges. Crucially, we are interested in graphs that satisfy semantic (edge-label) consistency criteria $\Omega \in \{0, 1\}^{c \times c}$, where $\Omega[c_i, c_j] = 0$ indicates edge types c_i and c_j can be adjacent, and 1 otherwise. A graph G satisfies Ω -constraint, i.e.:

$$P_{\Omega}(E) = \text{True}, \text{ iff } \sum_{v=1}^n \sum_{\substack{i,j \in \mathcal{N}(v) \\ i \neq j}} \left\| \left(\mathbf{E}[i, v] \otimes \mathbf{E}[v, j]^{\top} \right) \odot \Omega \right\|_F = 0 \quad (1)$$

where $\mathcal{N}(v)$ is the set of neighbors of node v , $\otimes$ denotes tensor product, and $\odot$ denotes Hadamard product. Additionally, any available structural constraints can also be incorporated, e.g., tree structure for airways.

Overview: We adopt a two-stage graph generation strategy [14], which is described in Fig. 1. In the first stage, a denoising diffusion probabilistic model [5] is trained on the node coordinates. Specifically, Gaussian noise is added to coordinates $\mathbf{X}^0$ through the forward process: $\mathbf{X}^t := \sqrt{\alpha_x^t} \mathbf{X}^{t-1} + \sqrt{1 - \alpha_x^t} \boldsymbol{\epsilon}$, $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, where α_x^t transitions from 1 to 0 with increasing t . A neural network f_θ predicts the noise $\hat{\boldsymbol{\epsilon}} := f_\theta(\mathbf{X}^t, t)$ using $\mathbf{X}^t := \sqrt{\bar{\alpha}_x^t} \mathbf{X}^0 + \sqrt{1 - \bar{\alpha}_x^t} \boldsymbol{\epsilon}$; ($\bar{\alpha}_x^t := \prod_{s=1}^t \alpha_x^s$). The model is trained using $\mathcal{L}_\theta := \mathbb{E}_{t, \mathbf{X}^0, \boldsymbol{\epsilon}} \|\boldsymbol{\epsilon} - \hat{\boldsymbol{\epsilon}}\|^2$. During sampling, node coordinates are denoised iteratively starting from $\mathbf{X}^T, \boldsymbol{\epsilon}_b \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$:

$$\mathbf{X}^{t-1} = \frac{1}{\sqrt{\alpha_x^t}} \left(\mathbf{X}^t - \frac{1 - \alpha_x^t}{\sqrt{1 - \bar{\alpha}_x^t}} \hat{\boldsymbol{\epsilon}}^t \right) + \sqrt{1 - \alpha_x^t} \boldsymbol{\epsilon}_b.$$

In the second stage, a discrete diffusion process [21, 14] is used on the edges, where a transition matrix $\mathbf{Q}^t := \alpha_e^t \mathbf{I} + (1 - \alpha_e^t) \mathbf{1}_e \mathbf{m}^\top$ changes the edges independently during forward noising as $\mathbf{E}^t := \mathbf{E}^{t-1} \mathbf{Q}^t$. A non-rotation-equivariant graph transformer g_ϕ is trained to predict $\hat{\mathbf{E}}^0 := g_\phi(\mathbf{E}^t, \mathbf{X}^0, t)$. The model is trained using cross-entropy loss using $\mathcal{L}_\phi := \mathbb{E}_{t, \mathbf{E}^0} \sum_{ij} CE(\mathbf{E}_{ij}^0, \hat{\mathbf{E}}_{ij}^0)$. During sampling, we sample edges from the posterior:

$$p_\phi(\mathbf{E}_{ij}^{t-1} | \mathbf{E}_{ij}^t) := \frac{\mathbf{E}_{ij}^t \mathbf{Q}^{t-1} \odot \hat{\mathbf{E}}_{ij}^0 \bar{\mathbf{Q}}_{ij}^{t-1}}{\hat{\mathbf{E}}_{ij}^0 \mathbf{Q}_{ij}^t \mathbf{E}_{ij}^{t-1}},$$

starting from $\mathbf{E}_{ij}^T \sim \mathbf{m}$, where $\bar{\mathbf{Q}}^t := \mathbf{Q}^1 \dots \mathbf{Q}^t$. Note that Prabhakar and Shit et al. [14] use uniform $\mathbf{m}$ and can not guarantee structural and semantic constraints; thus, it is prone to generating *semantically invalid* graphs.

Edge-deletion-based Noising: Discrete denoising diffusion models typically rely on uniform or marginal noise to perturb graph connectivity [21, 20, 14]. However, real-world spatial graphs are often sparse (i.e., the number of edges grows linearly with the number of nodes), so standard noise schedulers introduce excessive, implausible connections early on. This contradicts the assumption that samples should remain near the data manifold during early diffusion steps, leading to inefficiency. Recently, [8, 1] showed that discrete diffusion models benefit from an *absorbing state* noising process, mapping each token to a [MASK] in language modeling. Madeira et al. [9] adapted this to graphs by treating the absence of an edge as the [MASK] token, so each edge is either deleted or retains its existing label (never changing labels). The transition matrix for noising step is $\mathbf{Q}^t := \alpha_e^t \mathbf{I} + (1 - \alpha_e^t) \mathbf{1}_e \mathbf{e}_E^\top$, where $\mathbf{e}_E$ is the one hot encoding of the absence of edge. At $t = T$, all the edges are deleted, and we get an empty graph with only node coordinates. During sampling, only new edges are inserted, and once inserted, they remain fixed throughout denoising.

Semantically Consistent Projection: During edge denoising at time step t , $\hat{\mathbf{E}}^{t-1}$ is sampled from categorical distribution $p_\phi(\mathbf{E}^{t-1}|\mathbf{E}^t)$. Since the forward noising is designed by edge deletion, during denoising, the network learns to add edges. There is no guarantee that $\hat{\mathbf{E}}^{t-1}$ will satisfy structural or semantic constraints. Some newly added edges can violate structure and/or semantic constraints, creating an invalid intermediate graph. Recently, [9] showed that if the edge-deletion noising process preserves a structural constraint, then projecting intermediate graphs during denoising onto a constrained set ensures compliance with that constraint.

This key result motivates us to investigate semantic properties relevant to spatial biological graphs that remain invariant under the edge-deletion noising model. Since the edge-deletion process does not change the relationships among neighboring edges, we can devise a label-aware projector module to fix the intermediate graph if needed. Importantly, existing work [9] only considered structural properties, resulting in either accepting or rejecting new edges. This is useful to prevent the formation of cycles in many biological graphs, e.g., airways. However, in the case of multi-class edges, the model can propose a structurally consistent edge that only violates semantic constraints. A naive strategy of simply deleting these edges is suboptimal, as it penalizes the existence of an edge based on one sampled semantic label, resulting in overly sparse graphs.

To this end, we propose a novel label-aware projector, which allows us to sample from the posterior distribution constrained by our semantic plausibility. The projector begins with $E^{t-1} \leftarrow E^t$ and adds edges in E^{t-1} that satisfy the semantic-consistency criteria. Specifically, for each edge $e_{ij} \in \hat{E}^{t-1} \setminus E^t$ our projector adds $E^{t-1} \leftarrow E^{t-1} \cup \{e_{ij}\}$ if $P_\Omega(E^{t-1} \cup \{e_{ij}\}) = \text{True}$. If not, we sample k new edge labels from $e_{ij} \sim p_\phi(\mathbf{E}_{ij}^{t-1}|\mathbf{E}_{ij}^t)$. Note that this is in contrast with [9], which either accept or discard edges, whereas we harness stochastic variation in favor of increased new edge proposals. The projector accepts the edge if any of the k resampled edge labels satisfy the semantic consistency; otherwise, it rejects the edge. This strategy is visually demonstrated in Fig. 1. Such a strategy allows the model to fix invalid edges stochastically and avoids overly sparse graphs, leading to semantically valid outputs. Crucially, our semantic projector can be seamlessly combined with a structural projector.

4 Experiments

Datasets: We use two publicly available datasets for the experiments. First is the circle of Willis (CoW) dataset consisting of 300 MRA from CROWN [22] and 90 from TopCoW [24] challenges. Please note that we use only CROWN samples to train the diffusion model. The TopCoW samples contain verified segmentations and are used to evaluate the vessel labeler. The CROWN graphs generated following [14] contain 14 edge classes representing different artery labels. The second dataset is from the Airway Tree Modeling (ATM) [26] challenge, consisting of 300 CT images. We extracted the graph from the ground-truth segmentation using Voreen [2], then used radius information to define five thickness-based edge

classes (including a background/no-edge class). We manually removed erroneous samples, yielding a final training set of 251 graphs. The number of nodes ranges between 40 and 500 in ATM (average 127) and between 13 and 27 in CoW (average 8). Such diversity underscores the generalizability of our method.

Baselines: We use the two-stage diffusion model proposed by [14] as our baseline. We refer to it as Uniform based on its noising scheme. Additionally, we use Construct [9], which uses only structural constraints. We also include an ablation of our method, namely ours ($k=0$), that deletes edges with label violations while our model resamples k ($= 4$) times to fix these violations.

Metrics: Following [14,3,7], we evaluate by comparing various physical features of the generated samples, including node degree ($\deg(\mathcal{V})$), number of edges ($|\mathcal{E}|$), length of edges ($l_{\mathcal{E}}$), and angle between the edges ($\mathcal{E} \angle$). We also compute Betti-0 (β_0) for connected components and Betti-1 (β_1) for cycles. In addition, we report semantic validity (in %) for the generated graphs. We report balanced accuracy (bal. acc.) and macro averaged F1 (avg. F1) scores for the vessel labeler and link prediction tasks. Additionally, we report semantic validity (in %) for the latter.

Semantic Validity (S.V.): We use edge-label based criterion for semantic validity. **CoW:** The CoW graph has no structural constraints, and the TopCoW annotation protocol [24] determines its validity, please refer [24] for details. For CoW, we have 13×13 Ω matrix for 13 edge-labels. **ATM:** An airway tree is invalid if (i) it contains a cycle, and (ii) it violates radius-based label hierarchy (e.g., connects trachea with bronchioles). We have 4×4 Ω as shown in Fig. 1.

Downstream Tasks: We assess our model on two downstream tasks (Fig. 3). (i) **Vessel labeling** involves inferring multi-class labels to an unlabeled CoW graph based on the TopCoW annotation protocol [24]. We generate 1000 synthetic samples to train the labeler and evaluate it on *unseen* TopCoW samples. Specifically, we use a graph transformer to predict edge labels for the unlabeled edges and assess its performance against the reference labels. We use 30 samples as a test set and a 50/10 (train/val) split for the real-data baseline. (ii) **Link prediction** involves predicting plausible connections on an ATM graph with approximately 30% missing edges. We feed this incomplete graph into our diffusion model for 100 denoising steps. Because our model only adds edges during denoising, it can predict missing links. In contrast, [14] relies on a uniform noising process and thus cannot be used for link prediction. To the best of our knowledge, this is the first attempt to use an out-of-the-box diffusion model for link prediction on biological graphs.

Implementation Details: We use the same node and edge denoising architecture as proposed by [14]. We use 8 edge-denoising blocks for CROWN and 5 for ATM dataset. All models (including baselines) are trained from scratch for 1000 epochs on a single RTX A6000 GPU, with a batch size of 64 for CROWN and 4 for ATM with AdamW optimizer with a learning rate of $3e-4$. The discrete-diffusion model uses a linear noise schedule and 500 noising steps.

Table 1. Comparison of our method against other graph generation methods. Our proposed solution adheres to the data distribution statistics while satisfying the semantic validity (S.V.) constraints. KL statistics are reported in 10^{-3} and S.V. in %.

	Methods	$\deg(\mathcal{V}) \downarrow$	$ \mathcal{E} \downarrow$	$l_{\mathcal{E}} \downarrow$	$\mathcal{E} \angle \downarrow$	$\beta_0 \downarrow$	$\beta_1 \downarrow$	S.V.% $\uparrow$
CROWN	Uniform [14]	1.328	9.113	4.718	9.550	3.284	1.069	70
	Construct [9]	1.017	5.867	0.862	5.178	1.510	2.833	75
	Ours (k=0)	1.071	6.054	0.789	5.199	1.348	2.737	100
	Ours	1.014	4.893	0.864	5.174	1.047	2.007	100
ATM	Uniform [14]	8.432	2.696	5.074	4.369	3.077	1.425	10
	Construct [9]	2.301	2.734	6.089	4.350	9.065	0.0	40
	Ours (k=0)	2.319	3.707	6.112	4.349	10.31	0.0	100
	Ours	2.278	2.734	6.091	4.340	9.070	0.0	100

Table 2. Vessel labeling on *TopCoW*.

Train Data	bal. acc. $\uparrow$	avg. F1 $\uparrow$
Uniform [14]	88.44	0.874
Construct [9]	89.71	0.899
Ours (k=0)	94.68	0.961
Ours	95.45	0.967
Real data	93.77	0.946

Table 3. Link prediction on *ATM*. Note that [14] can not be used for link prediction, while our method can predict the missing links.

Method	bal. acc. $\uparrow$	avg. F1 $\uparrow$	S.V. $\uparrow$
Construct [9]	84.32	0.842	70
Ours (k=0)	84.83	0.849	100
Ours	84.97	0.849	100

Results on Circle of Willis: In Table 1, which compares the statistics of the generated samples, we note that [14] performs well in structural metrics, but generates only 70% semantically valid samples. Construct performs slightly better with 75% semantic validity. Please note that CROWN has no structural constraints. Hence, the improved results compared to [14] highlight the advantage of the edge-deletion noising scheme over the uniform scheme; Ours(k=0) uses semantic constraints, and thus, it generates 100% semantically-valid samples. However, deleting invalid edges makes the generated samples more sparse with worse distribution statistics. Our method leverages stochastic resampling and generates 100% valid samples with better distribution compliance. The effect is further exemplified by the superior performance of the downstream vessel labeling task trained on the synthetic graphs (c.f. Tab. 2 and Fig. 3).

Results on lung airway: Table 1 shows that the baseline method generates only 10% valid samples. Although it has the best Betti-0 performance compared to other methods, many connections are invalid and anatomically implausible (Fig. 2). Construct uses a tree constraint and generates 43 % valid samples. Ours (k=0) strategy produces semantically valid samples, but has worse distribution statistics. In contrast, our method produces 100% semantically valid samples with better adherence to the data distribution. Importantly, our model shows emerging link prediction capabilities on 30 test samples, which are not seen during training, as shown in Tab. 3 and Fig. 3.

Minimal Intervention During Sampling: In our experiments, the projection affects only 2% of all the generated edges. This limited intervention acts as




Fig. 2. Qualitative examples of graphs generated by [14], Construct, and our method compared to ground truth samples. The constraint violations are highlighted in red. [14] produces samples with label and structural violations (cycles in airway trees). Construct can produce samples with label violations. Our method tries to adhere to both label and structural constraints, potentially fixing the issues (shown in blue).

a safeguard against error accumulation during sampling, rather than consistently overwriting the model’s predictions. Hence, such minimal intervention sets our approach apart from largely post-hoc corrected random graph generators.

Limitations: While we propose the first semantically consistent biological graph generator, our work considers solely the edge-deletion invariants and not the edge-insertion one. Hence, it can not guarantee constraints such as generating a single connected component. Additionally, since our point cloud generator is fixed during edge denoising, we can not refine any node coordinates, thus limiting generation flexibility. We aim to address these limitations in future work.

5 Conclusion

Maintaining semantic validity is a significant obstacle to adopting diffusion-based biological graph generation methods. In this work, we ensure that generated graphs adhere to both structural and edge-label constraints. To the best of our knowledge, this is the first attempt to correct potential inconsistencies in generated samples stochastically. It makes our method particularly attractive to scenarios where a combination of multiple constraints reduces the set of acceptable outputs, rendering a naive rejection strategy suboptimal. Most importantly, our method addresses the major limitation of current data-driven biological graph generation models and makes discrete diffusion models more practical.




Fig. 3. Downstream applications of the diffusion model. **Left:** The vessel labeler receives circle of Willis graphs without any label and assigns multi-class labels based on TopCoW annotation protocol [24]. **Right:** The diffusion model predicts missing links on the airway tree samples. We show the ground truth samples for reference.

Acknowledgments. This work has been supported by the Helmut Horten Foundation. S. S. is supported by the UZH Postdoc Grant (K-74851-03-01). H. B. Li is supported by an SNF postdoctoral mobility grant.

Disclosure of Interests. The authors declare no competing interests for this article.

References

1. Austin, J., Johnson, D.D., Ho, J., Tarlow, D., Van Den Berg, R.: Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems* **34**, 17981–17993 (2021)
2. 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), 1–28 (2021)
3. Feldman, P., Fainstein, M., Siless, V., Delrieux, C., Iarussi, E.: Vesselvae: Recursive variational autoencoders for 3d blood vessel synthesis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 67–76. Springer (2023)
4. Gruver, N., Stanton, S., Frey, N., Rudner, T.G., Hotzel, I., Lafrance-Vanasse, J., Rajpal, A., Cho, K., Wilson, A.G.: Protein design with guided discrete diffusion. *Advances in neural information processing systems* **36**, 12489–12517 (2023)
5. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
6. Jessen, E., Steinbach, M.C., Debbaut, C., Schillinger, D.: Rigorous mathematical optimization of synthetic hepatic vascular trees. *Journal of the Royal Society Interface* **19**(191), 20220087 (2022)

7. Kuipers, T.P., Konduri, P.R., Marquering, H., Bekkers, E.J.: Generating cerebral vessel trees of acute ischemic stroke patients using conditional set-diffusion. In: *Medical Imaging with Deep Learning* (2024)
8. Lou, A., Meng, C., Ermon, S.: Discrete diffusion modeling by estimating the ratios of the data distribution. *arXiv preprint arXiv:2310.16834* (2023)
9. Madeira, M., Vignac, C., Thanou, D., Frossard, P.: Generative modelling of structurally constrained graphs. *Advances in Neural Information Processing Systems* **37**, 137218–137262 (2025)
10. Van de Moortele, T., Goerke, U., Wendt, C.H., Coletti, F.: Airway morphology and inspiratory flow features in the early stages of chronic obstructive pulmonary disease. *Clinical Biomechanics* **66**, 60–65 (2019)
11. Nousias, S., Zacharaki, E.I., Moustakas, K.: Avatree: An open-source computational modelling framework modelling anatomically valid airway tree conformations. *PloS one* **15**(4), e0230259 (2020)
12. 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)
13. Pothuau, L., Porion, P., Lespessailles, E., Benhamou, C., Levitz, P.: A new method for three-dimensional skeleton graph analysis of porous media: application to trabecular bone microarchitecture. *Journal of microscopy* **199**(2), 149–161 (2000)
14. Prabhakar, C., Shit, S., Musio, F., Yang, K., Amiranashvili, T., Paetzold, J.C., Li, H.B., Menze, B.: 3d vessel graph generation using denoising diffusion. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 3–13. Springer (2024)
15. Rahaghi, F.N., Ross, J.C., Agarwal, M., González, G., Come, C., Diaz, A., Vegas-Sánchez-Ferrero, G., Hunsaker, A., Estépar, R.S.J., Waxman, A., et al.: Pulmonary vascular morphology as an imaging biomarker in chronic thromboembolic pulmonary hypertension. *Pulmonary circulation* **6**(1), 70–81 (2016)
16. Rauch, N., Harders, M.: Interactive Synthesis of 3D Geometries of Blood Vessels. In: *Eurographics 2021 - Short Papers*. The Eurographics Association (2021). <https://doi.org/10.2312/egs.20211012>
17. Reichold, J., Stapanoni, M., Keller, A.L., Buck, A., Jenny, P., Weber, B.: Vascular graph model to simulate the cerebral blood flow in realistic vascular networks. *Journal of Cerebral Blood Flow & Metabolism* **29**(8), 1429–1443 (2009)
18. Savinkov, R., Grebennikov, D., Puchkova, D., Chereshev, V., Sazonov, I., Bocharov, G.: Graph theory for modeling and analysis of the human lymphatic system. *Mathematics* **8**(12), 2236 (2020)
19. Sharma, K., Kumar, S., Trivedi, R.: Diffuse, sample, project: plug-and-play controllable graph generation. In: *Forty-first International Conference on Machine Learning* (2024)
20. Vignac, C., Krawczuk, I., Siraudin, A., Wang, B., Cevher, V., Frossard, P.: Digress: Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734* (2022)
21. Vignac, C., Osman, N., Toni, L., Frossard, P.: Midi: Mixed graph and 3d denoising diffusion for molecule generation. *arXiv preprint arXiv:2302.09048* (2023)
22. Vos, I., Ruigrok, Y., Kuijff, H.: Results of the crown challenge on automated assessment of circle of willis morphology. In: *Medical Imaging with Deep Learning* (2024)

23. 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)
24. 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 preprint arXiv:2312.17670 (2024)
25. Yang, N., Zeng, K., Lu, H., Wu, Y., Yuan, Z., Chen, D., Jiang, S., Wu, J., Wang, Y., Yan, J.: Morphgrowing: A synchronized layer-by-layer growing approach for plausible neuronal morphology generation. arXiv preprint arXiv:2401.09500 (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)