

# 3D Vessel Graph Generation Using Denoising Diffusion

Chinmay Prabhakar<sup>\*1</sup> , Suprosanna Shit<sup>\*1</sup> , Fabio Musio<sup>1,2</sup>,  
Kaiyuan Yang<sup>1</sup>, Tamaz Amiranashvili<sup>1,3</sup> , Johannes C. Paetzold<sup>4</sup> ,  
Hongwei Bran Li<sup>5</sup> , and Bjoern Menze<sup>1</sup> 

<sup>1</sup> Department of Quantitative Biomedicine, University of Zurich, Switzerland

<sup>2</sup> Center for Computational Health, ZHAW, Switzerland

<sup>3</sup> Department of Computer Science, Technical University of Munich, Germany

<sup>4</sup> Department of Computing, Imperial College London, UK

<sup>5</sup> Athinoula A. Martinos Center, Harvard Medical School, USA

chinmay.prabhakar@uzh.ch

**Abstract.** Blood vessel networks, represented as 3D graphs, help predict disease biomarkers, simulate blood flow, and aid in synthetic image generation, relevant in both clinical and pre-clinical settings. However, generating realistic vessel graphs that correspond to an anatomy of interest is challenging. Previous methods aimed at generating vessel trees mostly in an autoregressive style and could not be applied to vessel graphs with cycles such as capillaries or specific anatomical structures such as the Circle of Willis. Addressing this gap, we introduce the first application of *denoising diffusion models* in 3D vessel graph generation. Our contributions include a novel, two-stage generation method that sequentially denoises node coordinates and edges. We experiment with two real-world vessel datasets, consisting of microscopic capillaries and major cerebral vessels, and demonstrate the generalizability of our method for producing diverse, novel, and anatomically plausible vessel graphs.

## 1 Introduction

Studying blood vessel networks as a 3D spatial graph offers a compact representation of anatomical and physiological properties of the circulatory system and, hence, has gained increasing interest in clinical [13] and pre-clinical image analysis [15]. Vessel graphs previously have been used to predict disease biomarkers [16], simulate blood flow [21], and create synthetic image-label pairs for segmentation [23, 14]. However, annotating large-scale samples is expensive [29] and tedious [23]. To mitigate this issue, synthetic vessel graphs can augment real datasets for downstream tasks. While extracting vessel graphs from a segmented volume has been studied numerous times [2, 15], generating realistic vessel graphs remains relatively underexplored. Previous methods [22, 14] primarily rely on rule-based vessel-tree generation based on a predefined oxygen

---

\* Contributed equally

concentration or nutrition distribution. However, such methods require precise domain modeling and often require tuning a substantial hyperparameter space. Therefore, a notable void exists, which is the generation of realistic 3D vessel graphs in an entirely data-driven way.

Recently, diffusion model-based graph generation has been successfully applied to various graph generation tasks ranging from molecule generation [24], protein design [30], and material sciences [28]. While these spatial graph generation tasks share a common goal, they differ significantly from each other in terms of the inductive bias describing the physical quantity of interest and how these quantities influence each other. For example, in the case of molecule generation, two neighboring atoms (nodes) being carbons provide a strong bias for the existence of a covalent bond. Such inductive bias in the form of node class, however, does not apply to vessel generation. Hence, the model has to learn entirely different co-dependencies and distributions; therefore, existing spatial graph generation methods are not readily applicable to vessel graph generation tasks. Furthermore, depending on the imaging resolutions, vessel graphs can have different topological and physiological properties. For example, capillaries are visible in microscopic images and contain a substantially higher number of cycles than large arteries, which are visible in magnetic resonance images of the brain. Hence, we aim to develop a flexible data-driven model capable of generating vessel graphs across different imaging resolutions.

In this work, we introduce the first denoising diffusion-based method for generating vessel graphs, addressing the challenge of simultaneously generating nodes and edges – an issue we identify as ill-posed for vessel graphs. Our approach focuses on first generating node coordinates, followed by the denoising of edges. This strategy allows the network to accurately model node distribution before learning to connect these nodes into a coherent vascular graph topology. Our experiments demonstrate the capability of our method to produce valid 3D vessel graphs ranging from capillary to large vessel levels using real vessel data.

To summarize, our contributions are threefold. Firstly, we introduce a novel diffusion denoising-based vessel graph generation method, marking a first in this domain. Secondly, we address the complexity of generating vessel graphs by proposing a tailored approach, sequentially denoising nodes and then edges. Lastly, through experimentation with two real-world vessel datasets – one representing capillaries and the other large vessels – we validate our method’s ability to generate diverse, unique, and valid vascular graphs.

## 2 Related Literature

**Vessel Graph Generation:** Previous methods have investigated autoregressive vessel tree generation methods. Schneider et al. [22] used underlying oxygen concentration to generate an arterial tree. However, such simulations are costly, and recently, Rauch et al. [20] proposed a computationally efficient statistical algorithm. The resultant graphs are useful for generating synthetic image-segmentation pairs to train segmentation models [10,14]. Parallely, there have

been attempts to generate vessel mesh. Wolternick et al. [27] proposed a generative adversarial network to synthesize coronary arteries. Recently, Feldman et al. [4] used a variational-auto-encoder to generate intracranial vessel segments. However, these network generation algorithms can not generate cycles in vessel graphs, which are crucial features for many anatomical scenarios such as capillary or Circle of Willis (CoW).


**Diffusion-based Graph Generation:** Recent advancements in diffusion models have led to the development of graph generation in various domains. Huang et al. [8] proposed a method for generating permutation-invariant graphs. Jo et al. [9] applied a score-based model for the graph generation process. Luo et al. [12] proposed a spectral denoising technique to increase the generation speed of graph models. Vignac et al. [24] developed a method for discrete denoising diffusion for categorical node and edge variables. Concurrently, Haefeli et al. [5] also utilize discrete diffusion models for graph generation. Peng et al. [17] tackled the issue of atom-bond inconsistencies in 3D molecule generation. Hua et al. [7] proposed a unified diffusion approach for molecule generation. Vignac et al. [25] merged discrete and continuous 3D denoising techniques to enhance molecule structure generation. Pinheiro et al. [19] proposed generating molecular structures within 3D spaces using voxel grids. Despite many works on molecular graph generation, there is a lack of diffusion-based vessel graph generation algorithms.

### 3 Methodology

We want to generate spatial graphs with nodes embedded in 3D space and edges with categorical attributes. Anatomically, the nodes represent the bifurcation points of the vascular network, and the edges represent the vessels connecting them. We denote  $\mathbf{x}_i \in \mathbb{R}^3$  for the 3D coordinates of the  $i^{\text{th}}$  node and  $\mathbf{e}_{ij} \in \mathbb{R}^c$  for the edge between the  $i^{\text{th}}$  and  $j^{\text{th}}$  node, where  $c$  is the number of edge classes. Thus, a graph is encoded by its node coordinates  $\mathbf{X} \in \mathbb{R}^{n \times 3}$  and categorical edge adjacency matrix  $\mathbf{E} \in \mathbb{R}^{n \times n \times c}$  where  $n$  is the number of nodes.

#### 3.1 Two-stage Denoising

In typical graph domains, e.g., for molecular graphs, the atom type and rotation equivariance provide a strong bias to determine the edge class. However, a node in a vascular graph does not exhibit the same property. The vascular nodes are often assigned edge-derived attributes, such as the node degree and vessel radius, and do not necessarily contain discriminative features for edge prediction. Thus, the node location is the principal predictor of the edge existence [26]. As a result, we believe that changing the node coordinates makes edge-type prediction difficult. Once all the node locations are known, search for plausible edge configuration benefits from the collective inductive bias of the node coordinates. In light of the above, we propose a two-stage graph generation strategy (c.f. Fig. 1). First, we focus on generating a plausible set of node coordinates as point clouds. In the second stage, we learn the edges, keeping the node coordinates fixed.



**Fig. 1.** Overview of our method. **Left.** First, we use a continuous diffusion model to generate plausible node locations. **Right.** Subsequently, we apply discrete diffusion to generate a plausible edge configuration given the node coordinates. The node coordinates remain unchanged during edge learning. The model is optimized by  $\mathcal{L}_{\text{Edge}}$  and focuses on the conditional edge distribution given node locations.

**Node Denoising:** We employ a denoising diffusion probabilistic model [6] for the node generation task. We add Gaussian noise to the node coordinates  $\mathbf{X}^0$  using a fixed noise model  $q(\mathbf{X}^t | \mathbf{X}^{t-1})$  as follows

$$\mathbf{X}^t := \sqrt{\alpha^t} \mathbf{X}^{t-1} + \sqrt{1 - \alpha^t} \boldsymbol{\epsilon}; \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \text{ and } t \in [1, T] \quad (1)$$

where  $\alpha^1, \alpha^2, \dots, \alpha^T$  controls the noise added to the coordinates. The scheduler is selected such that  $\alpha^1 \approx 1$  and  $\alpha^T \approx 0$ . This ensures that the coordinates have been mapped into a standard Gaussian.

Next, we denoise the node coordinates using a reverse diffusion process. During training, we select a random time point  $t$  and train the model to predict the added noise. We employ a lightweight neural net with few fully connected layers followed by multi-head cross-attention with time embedding. We denote the model  $f_\gamma$  with model parameter  $\gamma$ . We obtain  $\mathbf{X}^t := \sqrt{\bar{\alpha}^t} \mathbf{X}^0 + \sqrt{1 - \bar{\alpha}^t} \boldsymbol{\epsilon}$  and predict the added noise as  $\hat{\boldsymbol{\epsilon}} := f_\gamma(\mathbf{X}^t, t)$ , where  $\bar{\alpha}^t := \prod_{s=1}^t \alpha^s$ . We train  $f_\gamma$  by minimizing the following loss function.

$$\mathcal{L}_{\text{Node}} := \mathbb{E}_{t, \mathbf{X}^0, \boldsymbol{\epsilon}} \|\boldsymbol{\epsilon} - \hat{\boldsymbol{\epsilon}}\|^2 \quad (2)$$

**Edge Denoising:** Once the node denoising model is trained, we train an edge denoising model. For this, we adopt the discrete diffusion model [24, 25]. We start with edge attributes  $\mathbf{E}^0$  and employ the following noise model.

$$q(\mathbf{E}^t | \mathbf{E}^{t-1}) := \mathbf{E}^{t-1} \mathbf{Q}^t; \quad t \in [1, T] \quad (3)$$

where  $\mathbf{Q}_{ij}^t := q(e^t = j | e^{t-1} = i)$  is the Markov transition probability in the state-space of edge categories. For a marginal distribution of  $\mathbf{m}$ , the resultant transition matrix is constructed as  $\mathbf{Q}^t := \alpha^t \mathbf{I} + (1 - \alpha^t) \mathbf{1}_c \mathbf{m}'$ , where  $'$  indicates transpose operation and  $\alpha^t$  is the noise scheduler similar to the node diffusion.

Following Vignac et al. [25], we adopt the graph transformer network [3]  $f_\delta$  with parameter  $\delta$  to model our edge denoising network. Note that [25] is explicitly designed for molecular graphs and, uses rotation equivariant layers. Vascular graphs, on the other hand, can have a particular orientation based on the

anatomy of interest (e.g., in the case of the Circle of Willis). A rotation equivariant graph generation model can not capture such dataset property. Hence, our model is *not* rotation equivariant to preserve dataset orientation property. During training, we select a random time point  $t$  and obtain  $\mathbf{E}^t$  using Eq. 3 and predict the ground truth edge as  $\hat{\mathbf{E}}^0 := f_\delta(\mathbf{E}^t, \mathbf{X}^0, t)$ . The model is supervised by cross-entropy (CE) loss.

$$\mathcal{L}_{\text{CE}} := \mathbb{E}_{t, \mathbf{E}^0} \mathbb{E}_{q(\mathbf{E}^t | \mathbf{E}^0)} \sum_{ij} \text{CE}(\mathbf{E}_{ij}^0, \hat{\mathbf{E}}_{ij}^0) \quad (4)$$

In addition to the correct edge class, the node degree between edges is vital in vascular graphs [13]. Thus, we introduce a novel node degree loss. With multiple valid edge configurations for the given node locations, the model should predict a degree distribution akin to the ground truth. Thus, we use a KL divergence loss between prediction and target node degree distribution over a mini-batch. However, we need the adjacency matrix to compare the node degree, which requires discrete sampling from the predicted edge adjacency likelihood. This operation would break the gradient backpropagation. To address this issue, we use the Gumbel-softmax ( $GS$ ) trick for the adjacency matrix sampling. The node degree loss is computed as

$$\mathcal{L}_\circ := KL \left( p_{\text{deg}}(\mathbf{E}^0) || p_{\text{deg}}(GS(\hat{\mathbf{E}}^0)) \right) \quad (5)$$

The total loss for edge denoising is  $\mathcal{L}_{\text{Edge}} = \mathcal{L}_{\text{CE}} + \mathcal{L}_\circ$ . We study the contribution of these losses in our ablation study.

### 3.2 Graph Generation


Once both  $f_\gamma$  and  $f_\delta$  are trained, we can use it for sampling vessel graphs. First, we sample the number of nodes,  $n$ , from possible discrete values. Next, we sample  $\mathbf{X}^T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and perform the following denoising steps from  $t = T, T-1, \dots, 1$  to obtain the posterior  $p_\gamma(\mathbf{X}^{t-1} | \mathbf{X}^t)$  as follows:

$$\mathbf{X}^{t-1} = \frac{1}{\sqrt{\alpha^t}} \left( \mathbf{X}^t - \frac{1 - \alpha^t}{\sqrt{1 - \alpha^t}} \hat{\epsilon}^t \right) + \sqrt{1 - \alpha^t} \epsilon \quad (6)$$

where  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  if  $t > 1$ , else  $\epsilon = \mathbf{0}$ . Once we have  $\mathbf{X}^0$ , we sample  $\mathbf{E}_{ij}^T \sim \mathbf{m}$  to start the edge denoising steps from  $t = T, T-1, \dots, 1$ . We first obtain  $\hat{\mathbf{E}}^0$  from  $\mathbf{E}^t$  and compute the posterior distribution of the edge category using the following as proposed in [1, 24]

$$p_\delta(\mathbf{E}_{ij}^{t-1} | \mathbf{E}_{ij}^t) = \frac{\mathbf{E}_{ij}^t \mathbf{Q}_{ij}^{t'} \odot \hat{\mathbf{E}}_{ij}^0 \bar{\mathbf{Q}}_{ij}^{t-1}}{\hat{\mathbf{E}}_{ij}^0 \bar{\mathbf{Q}}_{ij}^t \mathbf{E}_{ij}^{t'}} \quad (7)$$

where  $\bar{\mathbf{Q}}^t := \mathbf{Q}^1 \dots \mathbf{Q}^t$  and  $\odot$  denotes pointwise multiplication. In the next step, we sample the categorical value of edge attributes from  $\mathbf{E}_{ij}^{t-1} \sim p_\delta(\mathbf{E}_{ij}^{t-1} | \mathbf{E}_{ij}^t)$ .



**Fig. 2.** Two real-world vessel datasets. **Left.** An example from VesSAP and the capillary level vessel graph patch. **Right.** The extracted metric graph for the blood vessel in CoW. (Best viewed when zoomed in)

## 4 Experiments

**Datasets:** We have experimented on two real-world publicly available datasets. The first one consists of microscopic images of capillary-level vessel graphs, namely VesSAP [23,15]. We have chosen 24 annotated volumes, and the extracted graphs from Voreen [2]. For our experiments, we used a patch size graph of  $48 \times 48 \times 48$  in the voxel space in the original resolution. We have used the radius information to create 4 edge classes based on thickness, including one background (no-edge) class. For the second dataset, we have the Circle of Willis (CoW) dataset consisting of 390 MRA images from CROWN<sup>6</sup> ( $n=300$ ) and TopCoW ( $n=90$ ) [29] challenges. We used a multi-class segmentation tool trained on the TopCoW dataset to segment the CROWN datasets and Voreen to extract graphs from the segmentation. For CoW data, we used the whole CoW graph as one sample, and it has 14 classes of edges containing different artery labels and a background class. For both datasets, we have used the refined graph [2].

**Baselines:** There is a notable void in data-driven vessel graph generation methods. Although tree-generation methods exist, they are not applicable in our setting since we are also interested in generating cycles in the graph. Therefore, we select two recent molecular graph generation methods, namely Congress [24] and MiDi [25], that can generate 3D spatial graphs with cycles as baselines.

**Metrics:** We acknowledge that finding a single good metric to benchmark vessel graph generation methods combining both the node coordinates and the edge information is difficult. It has been reported that when the generated graphs accurately represent real data distribution, they improve downstream performance [10]. Previous work [4] has used distributional differences in radius, total length, and tortuosity between generated and real mesh. Along this line, we adopt the following physical parameters to compute the KL divergence between real and generated graphs. For nodes, we use 3D coordinate positions ( $\{x, y, z\}$ ) and node degree ( $deg(\mathcal{V})$ ). For edges, we use total edge number ( $|\mathcal{E}|$ ), edge length ( $l_{\mathcal{E}}$ ), angle between edges ( $\mathcal{E} \angle$ ), edge orientation with three axes  $\{\theta, \phi, \psi\}$ , Bett-0 for connected component ( $\beta_0$ ), and Betti-1 for cycles ( $\beta_1$ ).

<sup>6</sup> <https://crown.isi.uu.nl/>


**Table 1.** Comparison of our method against the existing spatial graph generation methods. Our proposed two-stage solution outperforms the existing baselines, achieving the lowest KL divergence on both the VesSAP and the CoW datasets.

|        | Methods       | $x, y, z$    | $\deg(\mathcal{V})$ | $ \mathcal{E} $ | $l_{\mathcal{E}}$ | $\mathcal{E} \angle$ | $\theta, \phi, \psi$ | $\beta_0$    | $\beta_1$    |
|--------|---------------|--------------|---------------------|-----------------|-------------------|----------------------|----------------------|--------------|--------------|
| VesSAP | Congress [24] | 0.003        | 0.211               | 0.288           | 0.112             | 0.202                | 0.002                | 6.226        | 8.251        |
|        | MiDi [25]     | 0.011        | 0.023               | 0.002           | <b>0.003</b>      | 0.011                | 0.003                | 0.014        | 0.074        |
|        | <b>Ours</b>   | <b>0.002</b> | <b>0.006</b>        | <b>0.001</b>    | 0.006             | <b>0.005</b>         | <b>0.001</b>         | <b>0.001</b> | <b>0.068</b> |
| CoW    | Congress [24] | 0.003        | 0.010               | 0.046           | 0.065             | 0.081                | 0.005                | 0.051        | 0.032        |
|        | MiDi [25]     | 0.010        | 0.011               | 0.040           | 0.012             | <b>0.007</b>         | 0.005                | 0.049        | 0.028        |
|        | <b>Ours</b>   | <b>0.002</b> | <b>0.003</b>        | <b>0.004</b>    | <b>0.006</b>      | <b>0.007</b>         | <b>0.002</b>         | <b>0.015</b> | <b>0.017</b> |

**Implementation Details:** The node-denoising network works with normalized coordinates and uses two blocks, each containing 1) A multi-layer perception (MLP) to process the coordinates, 2) an MLP to process the denoising timestep (encoded as a sinusoidal positional embedding), and 3) A multi-head self-attention (MSA) processes the summed output of the two MLPs. A final MLP projects the features to  $\mathbb{R}^3$ . All the layers use a hidden dimension of size 256. The edge-denoising network is composed of eight blocks. Each block gets the node coordinates, noisy edges, and the timestep information as input. A feature transformation layer (FiLM) [18] processes the edges and timestep information and merges them with node features. The resultant node features are passed through an MSA. The final node features predict the output edges. Please refer to the supplementary for all models’ hyperparameters. All models (including baselines) are trained from scratch for 1000 epochs on a single A-6000 GPU with a batch size 64. We use AdamW [11] optimizer with a learning rate of 0.0003. We implement the code using the pyTorch and pyTorchGeometric libraries. Our code is available at [https://github.com/chinmay5/vessel\\_diffuse](https://github.com/chinmay5/vessel_diffuse)

**Results and Discussion:** Table 1 compares the statistics of the generated 3D vessel graphs with the state-of-art methods, viz., the discrete graph generation model MiDi and the continuous model, Congress. Congress yields overconnected graphs, failing to reflect the actual distribution. MiDi performs better but struggles with connectivity (node degrees,  $\beta_0$ ) and edge properties (length, number of edges). Our graph generation model outperforms the baseline methods on both CoW and VesSAP datasets, producing graphs that most closely emulate the sample statistics and obtain the lowest KL divergence for various graph and topology metrics. This result validates the utility of our two-stage design principle, which starts with node diffusion and then moves on to edge diffusion. Please refer to the supplementary for a detailed comparison of distribution.

Figure 3 shows visual examples of the graphs generated by our model and the ground truth. Our model is able to capture the complex vascular shapes in the VesSAP dataset and can generate cycles of similar complexity. The CoW dataset has a specific orientation of edges, which is again captured faithfully by our model. Please refer to the supplementary for more qualitative results.



**Fig. 3.** VesSAP and CoW graphs generated by our model in comparison to ground truth samples. Our model is able to learn complex structures, such as loops and orientation characteristics for both datasets.

**Table 2.** Ablation of rotation equivariance (R.E.), two-stage denoising model (II S), and the degree loss  $\mathcal{L}_o$  for CoW dataset. The best model is not rotation equivariant and uses two-stage denoising and degree loss. All values are reported in units of  $10^{-2}$ .

| R.E. | II S | $\mathcal{L}_o$ | $x, y, z$   | $\deg(V)$   | $ \mathcal{E} $ | $l_{\mathcal{E}}$ | $\mathcal{E} \angle$ | $\theta, \phi, \psi$ | $\beta_0$   | $\beta_1$   |
|------|------|-----------------|-------------|-------------|-----------------|-------------------|----------------------|----------------------|-------------|-------------|
| ✓    | ✗    | ✗               | 1.04        | 1.12        | 4.01            | 1.29              | 0.78                 | 0.52                 | 4.89        | 2.01        |
| ✗    | ✗    | ✗               | 0.44        | 1.89        | 4.43            | 6.08              | 8.14                 | 0.33                 | 11.16       | 4.31        |
| ✗    | ✓    | ✗               | <b>0.21</b> | <b>0.29</b> | 3.43            | 0.74              | 0.81                 | 0.19                 | <b>1.10</b> | 2.31        |
| ✗    | ✓    | ✓               | <b>0.21</b> | <b>0.29</b> | <b>0.40</b>     | <b>0.63</b>       | <b>0.74</b>          | <b>0.19</b>          | 1.54        | <b>1.71</b> |

**Ablation:** We analyze the importance of different components of our proposed method. We choose the Circle of Willis dataset for our ablation experiments. We summarize the results in Table 2. Disabling the rotation equivariance (R.E.) in the edge denoising blocks enables the model (Row 2) to learn about the vascular graphs’ inherent orientation bias. Although this leads to graphs with better orientation  $(\theta, \phi, \psi)$ , R.E. is a powerful inductive bias for edge connectivity, and naively turning it off leads to very sparse graphs (significant deterioration in  $\beta_0$ ). We switch to our two-stage (II S) denoising approach to counteract the increase in sparsity. We observe an immediate performance improvement, with most statistics better than the baseline (Row 3 *vs.* Row 1). We further improve the quality of the generated samples by incorporating a loss on the node degree distribution. This loss provides additional supervision for generating graphs and

further improves topological properties, leading to samples that most closely emulate the ground truth distributions (Row 4).

## 5 Conclusion

Generating spatial vascular graphs is a clinically relevant, albeit challenging task. In this work, we propose a two-stage approach to generate graphs that better capture the vascular intricacies. To the best of our knowledge, this is the first application of a diffusion model to generate vascular graphs. Since the diffusion process is at the heart of our method, we can easily extend our method to generate graphs conditioned on, for instance, disease labels. We defer such extensions to future work. Most importantly, diffusion models have revolutionized generative modeling, outperforming other competitive methods, and we hope our method paves the way for more widespread adoption of diffusion models for vascular graph generation.

**Acknowledgments.** This work has been supported by the Helmut Horten Foundation. S. S. is supported by the GRC Travel Grant from UZH. F.M. is funded by the DIZH grant. H. B. Li is supported by an SNF postdoctoral mobility grant.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of 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. Dwivedi, V.P., Bresson, X.: A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699* (2020)
4. 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)
5. Haefeli, K.K., Martinkus, K., Perraudin, N., Wattenhofer, R.: Diffusion models for graphs benefit from discrete state spaces. *arXiv preprint arXiv:2210.01549* (2022)
6. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
7. Hua, C., Luan, S., Xu, M., Ying, R., Fu, J., Ermon, S., Precup, D.: Mudiff: Unified diffusion for complete molecule generation. *arXiv preprint arXiv:2304.14621* (2023)
8. Huang, H., Sun, L., Du, B., Fu, Y., Lv, W.: Graphgdp: Generative diffusion processes for permutation invariant graph generation. In: *2022 IEEE International Conference on Data Mining (ICDM)*. pp. 201–210. IEEE (2022)

9. Jo, J., Lee, S., Hwang, S.J.: Score-based generative modeling of graphs via the system of stochastic differential equations. In: International Conference on Machine Learning. pp. 10362–10383. PMLR (2022)
10. Kreitner, L., Paetzold, J.C., Rauch, N., Chen, C., Hagag, A.M., Fayed, A.E., Sivaprasad, S., Rausch, S., Weichsel, J., Menze, B.H., et al.: Synthetic optical coherence tomography angiographs for detailed retinal vessel segmentation without human annotations. *IEEE Transactions on Medical Imaging* (2024)
11. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
12. Luo, T., Mo, Z., Pan, S.J.: Fast graph generative model via spectral diffusion. arXiv preprint arXiv:2211.08892 (2022)
13. Lyu, X., Cheng, L., Zhang, S.: The retina benchmark for retinal vascular tree analysis. *Scientific Data* **9**(1), 397 (2022)
14. 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)
15. 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)
16. Paetzold, J.C., Lux, L., Kreitner, L., Ezhov, I., Shit, S., Lotery, A.J., Menten, M.J., Rueckert, D.: Geometric deep learning for disease classification in octa images. *Investigative Ophthalmology & Visual Science* **64**(8), 1098–1098 (2023)
17. Peng, X., Guan, J., Liu, Q., Ma, J.: Moldiff: Addressing the atom-bond inconsistency problem in 3d molecule diffusion generation. arXiv preprint arXiv:2305.07508 (2023)
18. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
19. Pinheiro, P.O., Rackers, J., Kleinhenz, J., Maser, M., Mahmood, O., Watkins, A., Ra, S., Sresht, V., Saremi, S.: 3d molecule generation by denoising voxel grids. *Advances in Neural Information Processing Systems* **36** (2024)
20. 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>
21. Reichold, J., Stampanoni, 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)
22. Schneider, M., Reichold, J., Weber, B., Székely, G., Hirsch, S.: Tissue metabolism driven arterial tree generation. *Medical image analysis* **16**(7), 1397–1414 (2012)
23. Todorov, M.I., Paetzold, J.C., Schoppe, O., Tetteh, G., Shit, S., Efremov, V., Todorov-Völgyi, K., Düring, M., Dichgans, M., Piraud, M., et al.: Machine learning analysis of whole mouse brain vasculature. *Nature methods* **17**(4), 442–449 (2020)
24. 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)
25. Vignac, C., Osman, N., Toni, L., Frossard, P.: Midi: Mixed graph and 3d denoising diffusion for molecule generation. arXiv preprint arXiv:2302.09048 (2023)

26. Wittmann, B., Paetzold, J.C., Prabhakar, C., Rueckert, D., Menze, B.: Link prediction for flow-driven spatial networks. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2472–2481 (2024)
27. Wolterink, J.M., Leiner, T., Isgum, I.: Blood vessel geometry synthesis using generative adversarial networks. arXiv preprint arXiv:1804.04381 (2018)
28. Xie, T., Fu, X., Ganea, O.E., Barzilay, R., Jaakkola, T.: Crystal diffusion variational autoencoder for periodic material generation. arXiv preprint arXiv:2110.06197 (2021)
29. 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 (2023)
30. Yi, K., Zhou, B., Shen, Y., Liò, P., Wang, Y.: Graph denoising diffusion for inverse protein folding. *Advances in Neural Information Processing Systems* **36** (2024)