

Representing Anatomical Trees by Denoising Diffusion of Implicit Neural Fields

Ashish Sinha and Ghassan Hamarneh

Simon Fraser University
{ashish.sinha,hamarneh}@sfu.ca

Abstract. Anatomical trees play a central role in clinical diagnosis and treatment planning. However, accurately representing anatomical trees is challenging due to their varying and complex topology and geometry. Traditional methods for representing tree structures, captured using medical imaging, while invaluable for visualizing vascular and bronchial networks, exhibit drawbacks in terms of limited resolution, flexibility, and efficiency. Recently, implicit neural representations (INRs) have emerged as a powerful tool for representing shapes accurately and efficiently. We propose a novel approach for representing anatomical trees using INR, while also capturing the distribution of a set of trees via denoising diffusion in the space of INRs. We accurately capture the intricate geometries and topologies of anatomical trees at any desired resolution. Through extensive qualitative and quantitative evaluation, we demonstrate high-fidelity tree reconstruction with arbitrary resolution yet compact storage, and versatility across anatomical sites and tree complexities. The code will be available here¹.

Keywords: Anatomical Tree · Vasculature · Bronchial Tree · Representation Learning · Neural Fields · Implicit Neural Representations · Denoising Diffusion.

1 Introduction

Motivation for Studying Anatomical Trees. Accurate extraction and analysis of anatomical trees, including vasculature and airways, is crucial for predicting fatal diseases like ischemic heart disease [36] and aiding clinical tasks like surgical planning [30], computational fluid dynamics (CFD) [34], and disease prognosis [1]. However, the application of extracted trees for such purposes may be hindered by their explicit representation as meshes or volumetric grids.

Previous Works on Tree Representation. Prior works have utilized discrete and explicit representations, such as medial axis [25], minimal paths [14], grammar [15], voxel grids [42], and meshes [38], to represent anatomical trees. However, all these explicit representations have their own peculiar shortcomings,

¹ <https://github.com/sinashish/TreeDiffusion>

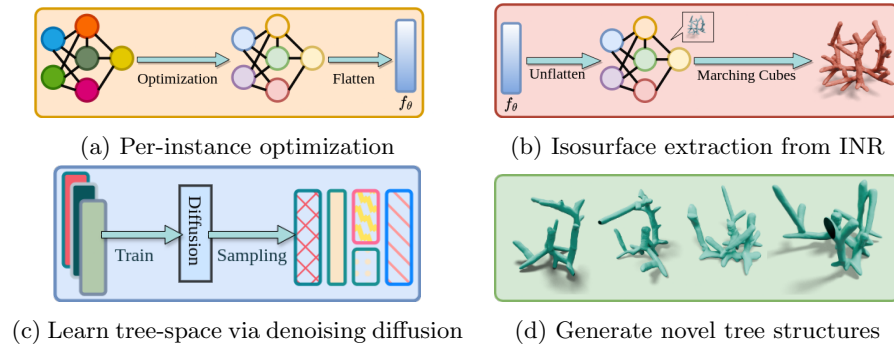


Fig. 1: **Overview:** (a) An INR is optimized for each sample in the training dataset, and then flattened to a 1D vector. (b) During inference, the INR is recovered from the flattened vector, followed by MC to extract the underlying signal. (c) The diffusion transformer takes the flattened vectors as input to model the diffusion process. After training, novel INRs can be sampled and used for downstream tasks via (b). (d) Novel tree structures visualized as mesh.

requiring complex ad-hoc machinery [12,28] to obtain a smooth surface essential for vascular modeling [1], and/or large memory footprint [19], up to $O(n^3)$ for voxels, to preserve fine details. Although rule-based representations, such as L-system, offer some degree of control over the produced structures [27], the complexity of rules and parameters and the computational cost of generating and rendering the shapes limit their use in current medical applications. Therefore, there is a need for a continuous shape representation that is independent of spatial resolution, is memory efficient, and can easily be integrated into deep learning pipelines.

Implicit Neural Representation (INR). Recently, INRs (or neural fields, implicit fields) have been proposed to address the shortcomings of explicit representations, which use multi-layer perceptrons (MLP) to fit a continuous function that *implicitly* represents a signal of interest [21,23,29] as level sets. This approach not only learns a representation of arbitrary shapes with a small number of parameters while still achieving a high surface accuracy, but also supports reconstructions at any arbitrary resolution. To this end, we propose to use INRs for representing anatomical trees and demonstrate their efficiency and effectiveness in modeling simple and complex, vasculature and airways, in 2D and 3D, on both synthetic and real data.

Generative Modeling of Tree Distributions. While implicit representations and level sets have been widely used for segmentation [2,18] and 3D modeling [11,13] of vessel structures, there are limited works that address the modeling of the “shape space” of tree structures. Previous works that attempt to analyze and synthesize tree structures are either *model-based* or *data-driven*.

Model-based approaches [8,9,33,39] record the tree topology along with branch attributes [7,41], *e.g.*, length and angle. However, for trees with diverse topologies (*e.g.*, different number of branches), a fixed-size representation no longer suffices, and altering the model size across trees impairs statistical analysis. *Data-driven* methods use generative modeling to represent the distribution of training data and sample the distribution to synthesize novel instances. GANs and VAEs were used to model the distribution of vessels [6,37] but did not leverage the descriptive power of INRs as [1,23] do. Diffusion models [10] were employed in [4] to learn the distribution on latent VAE codes of shapes but required large training sets for faithful encoding. For a more faithful representation, INRs were used in [5], though neither [4,5] modeled tubular structures or trees.

Our contributions To address the shortcomings of previous works, we are the first to make the following contributions: (1) We train INRs to achieve a faithful representation of complex anatomical trees and demonstrate their usage in segmenting trees from medical images; (2) we employ diffusion models on the space of INR-represented trees for learning tree distributions and generating plausible novel trees with complex topology; (3) we showcase the versatility of our approach in representing trees of varying dimensionality, complexity, and anatomy; and (4) We quantitatively assess our method’s compactness representation and reconstruction accuracy, at arbitrarily high resolution.

2 Methodology

Overview. We propose a diffusion-based generative approach comprising two steps. First, we optimize an INR realized as an MLP by overfitting its parameters to each tree in a dataset so that each INR faithfully represents each tree as a neural occupancy field (Fig. 1a). Second, the optimized INRs are flattened into 1D vectors and used to train a denoising diffusion model (DDM) that captures the statistical distribution of the trees (Fig. 1c). The DDM is used to synthesize new INRs, via a reverse diffusion process (Fig. 1b), which represents the occupancy field (or a mesh, if marching cubes [17] is applied) of the tree (Fig. 1d).

Stage 1: Learning per-sample INR. We represent each tree, in a dataset $\{S_i\}_{i=1}^N$ of N trees, as a neural occupancy field, or INR, $f(x; \theta) : \mathbb{R}^3 \rightarrow [0, 1]$, modelled as an MLP and parameterized by $\theta \in \mathbb{R}^D$. Note that, unlike prior works [23,40] that involve parameter sharing across the entire dataset, we optimize an INR separately for each tree sample.

Given $o_i(x)$, the ground truth occupancy of x for S_i , we optimize θ_i as:

$$\arg \min_{\theta_i} \sum_{x_j \in \mathbb{R}^3} \|f_{\theta_i}(x_j) - o_i(x_j)\|_2. \quad (1)$$

Similar to [21,22], our MLP architecture is a fully-connected network with L layers, of hidden size D , and ReLU activation functions. To learn high-fidelity

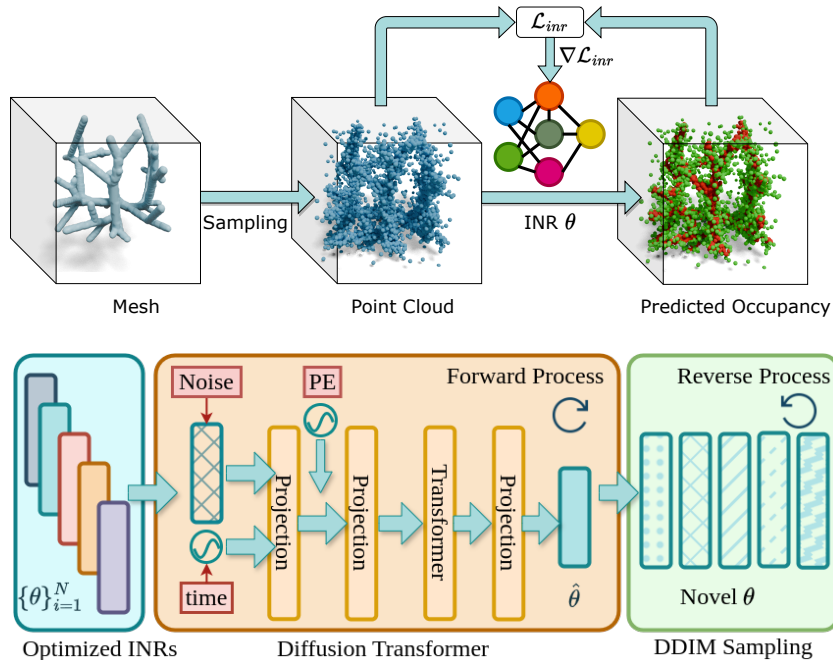


Fig. 2: **Pipeline:** Top: Given a 3D mesh, sampled points and GT occupancies (inside/outside), an INR (θ) is optimized to fit to the shape. Bottom: Optimized INRs are flattened to a 1D vector, fed to the transformer-based diffusion model $\mathcal{D}(\phi)$ and optimized to predict noise. The novel INRs can then be sampled from the trained $\mathcal{D}(\phi)$ in the reverse process.

neural representations of anatomical trees, we follow [21] to first instantiate all training instances S_i in Lipschitz domain $\Omega \in [-1, 1]$ and then adaptively sample points and ground-truth occupancies both inside and outside the 3D surface of S_i , to supervise the optimization of each MLP until convergence. Figure 2 illustrates this optimization process.

Stage 2: Diffusion on INRs. We train a diffusion model on the space of INRs, *i.e.*, the weights and biases $\{\theta_i\}_{i=1}^N$. Our transformer-based [35] diffusion model $\mathcal{D}(\phi)$ is inspired by [5,24], and takes the flattened 1D vectors of θ_i as input. However, before passing the input to $\mathcal{D}(\phi)$, each θ_i undergoes a *layer-by-layer* decomposition into k tokens by MLPs [24]. This is essential due to the potential variation in the dimensionality of different layers in θ_i and ensures correspondence of any element j across all trees $\theta_i, \forall i, j$. During forward diffusion, we apply noise η_t at timestep t to each vector θ_i to obtain a noisy vector θ_i^* .

Following [5,24], the sinusoidal embedding of t along with θ_i^* is fed to a linear projection layer, whose output is concatenated with a learnable positional encoding vector to obtain \mathcal{G} . The resultant \mathcal{G} is then passed to the transformer,

which outputs denoised INRs $\hat{\theta}_i$. Parameters ϕ of \mathcal{D} are optimized using:

$$\arg \min_{\phi} \frac{1}{N} \sum_{i=1}^N \|\hat{\theta}_i - \theta_i\|_2. \quad (2)$$

After training, we employ DDIM [31] sampling, *i.e.*, a reverse diffusion process, to sample new INRs by feeding-in noise as $\hat{\theta}_i$ to $\mathcal{D}(\phi)$, and gradually denoising it. Section 1 shows an illustration of the overall diffusion process.

3 Datasets, Results, and Implementation Details

Datasets. We use the following datasets to evaluate our approach: (1) VascuSynth: 120 synthetic 3D vascular trees, ranging from 1 branch to complex trees with multiple branches and bifurcations [9]; (2) IntRA: 103 3D meshes of intracranial vasculature extracted from MRA [38]; (3) BraTS: 40 3D MRI brain scans [20]; (4) HaN-Seg: 20 segmentations of Head&Neck CTA scans [26]; (5) DRIVE: 20 2D retinal fundus color images with vessel segmentation masks [32]; (6) EXACT: One CT scan of a bronchial tree [16]; (7) WBMRA: One whole body MRA; and (8) CoW: One circle of Willis mesh. Our evaluation covers the following aspects.

Fidelity and Compactness. Following [21], we use Chamfer distance (CD) to evaluate the fidelity of INRs in representing anatomical trees and summarize it in Tables 1 and 2. Figure 3 shows that we can achieve higher reconstruction accuracy with a smaller memory footprint. For example, as seen in Table 2 we need 68 MB for volumes and 12 MB for meshes compared to INRs that only occupy 0.75 MB and 0.63 MB, respectively, offering $\approx 90\times$ and $\approx 19\times$ compression, respectively, with minimal loss of reconstruction accuracy. Similarly, Table 1 shows that $10k$ parameters of INR are enough to encode the geometry of a 128^3 volume of IntRA, resulting in $\approx 220\times$ compression.

Table 1: Quantitative comparison of reconstruction accuracy and compression ratio on the VascuSynth (**V**) [9] and IntRA (**I**) [38] for different configurations of our MLP architecture. CD denotes chamfer distance (10^{-3}) between the reconstructed mesh and the ground truth, whereas compression denotes the ratio between size of INR and size of the original input modality, *i.e.*, raw volume. \uparrow denotes higher is better and \downarrow lower is better.

D	#Params (M)	Size _{inr} (MB)	L = 1				L = 3				L = 5							
			Compression (\uparrow)		CD (\downarrow)		Compression (\uparrow)		CD (\downarrow)		Compression (\uparrow)		CD (\downarrow)					
			V	I	V	I	V	I	V	I	V	I	V	I				
64	0.01	0.03	27.05	223.57	12.95	29.84	0.03	0.10	9.47	78.25	8.65	15.48	0.04	0.16	5.73	47.37	8.74	21.90
128	0.03	0.13	7.32	60.50	10.36	24.21	0.10	0.38	2.49	20.55	7.09	8.33	0.17	0.63	1.50	12.37	7.99	17.35
256	0.13	0.51	1.88	15.53	9.24	18.19	0.40	1.51	0.61	5.06	6.46	7.48	0.66	2.51	0.38	3.12	7.07	11.33
512	0.53	2.01	0.47	3.86	10.24	24.66	1.58	6.02	0.08	1.29	7.40	19.99	2.63	10.03	0.09	0.77	6.96	12.50
1024	2.10	8.03	0.04	0.96	11.37	40.06	6.30	24.04	0.04	0.32	7.37	16.58	10.50	42.48	0.02	0.19	7.64	17.26

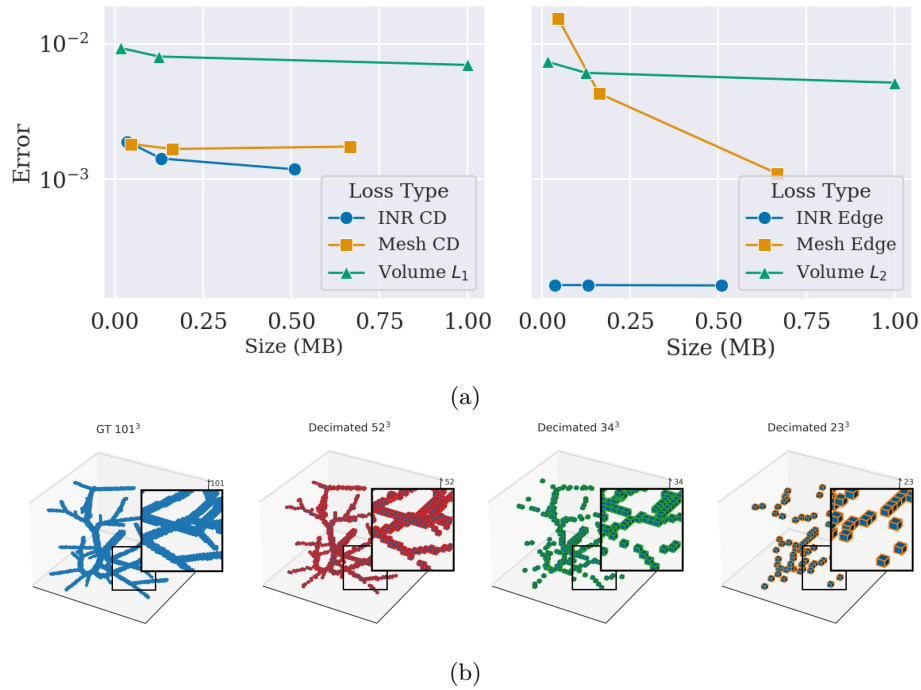



Fig. 3: **Compression vs Reconstruction Accuracy:** (a) We decimate/downscale the mesh/volume to occupy \approx same memory space as INR, and report the reconstruction error using Chamfer distance (CD) and edge length loss (Edge) for meshes and INRs, and L_1 and L_2 for volumes. Note the higher error for meshes and volumes w.r.t INRs for the same storage. (b) Illustration of GT and downsampled volumes. Notice the disconnected components.

Versatility. We assess the effectiveness of INRs in representing anatomical trees from various medical image modalities. Figure 4a and Table 2 qualitatively and quantitatively show the adaptability of INRs in modeling DRIVE vessel masks, 3D meshes, CTA and MRA volumes from VascuSynth, Intra, EXACT and CoW. Figure 4b depicts how the same MLP architecture can represent trees of varying complexity.

Arbitrary Resolution. Achieving high-quality mesh and volume reconstruction is crucial for the accurate representation of vasculature, *e.g.*, for fluid flow simulation. Figure 5 illustrates how INRs can reconstruct anatomical trees from both Intra and VascuSynth at various resolutions without the need for adjusting the model size or retraining.

Space of INR-based Trees. Once trees are represented via INRs, we study the space of INRs. The t-SNE plot in Figure 6b illustrates how trees with similar



(a) Various anatomical sites and imaging modalities (b) Various complexities

Fig. 4: **Versatility**: Visualization of different synthetic and real anatomical trees represented as an INR from various medical imaging modalities and organs. We normalize all shapes to $[-1, 1]$ and images to $[0, 1]$ to report the reconstruction error using MSE and CD ($\times 10^{-3}$) between **ground truth** and the underlying signal extracted from **INR**.

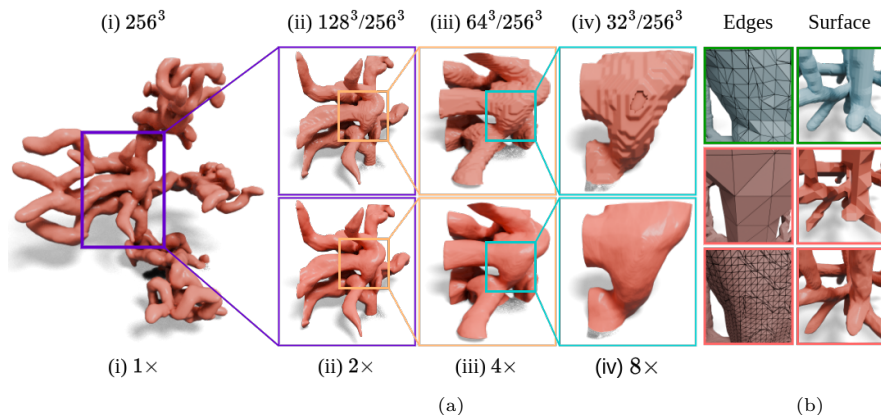


Fig. 5: **Arbitrary Resolution**: (a) Comparison of 2x, 4x, and 8x zoom on an InTRA sample represented as a volumetric grid (top) and INR (bottom). The resolution for each sub-figure is shown on top as $\text{volume}^3/\text{INR}^3$. Notice the smoothness of the surface even at 8x zoom. (b) Zoomed-in regions of a VascuSynth mesh **reconstructed from INRs** and **ground truth** at different mesh resolutions displayed using faces and edges.

number of bifurcations, *i.e.* topological complexities, tend to cluster in the 2D t-SNE space. However, we observe cases where trees with the same number of branches are distant, which may be due to how t-SNE projects high-dimensional samples to 2D thus distort inter-sample distances. Therefore, to further investigate, we compute the matrix of L_2 distances between INRs of VascuSynth trees with i and j bifurcations, where $i, j \in \{1, \dots, 12\}$ (Fig. 6a). The reported distances are the average across 10 trees per bifurcation count. We notice low (≈ 0) values along the diagonal, indicating that trees with the same branch count have

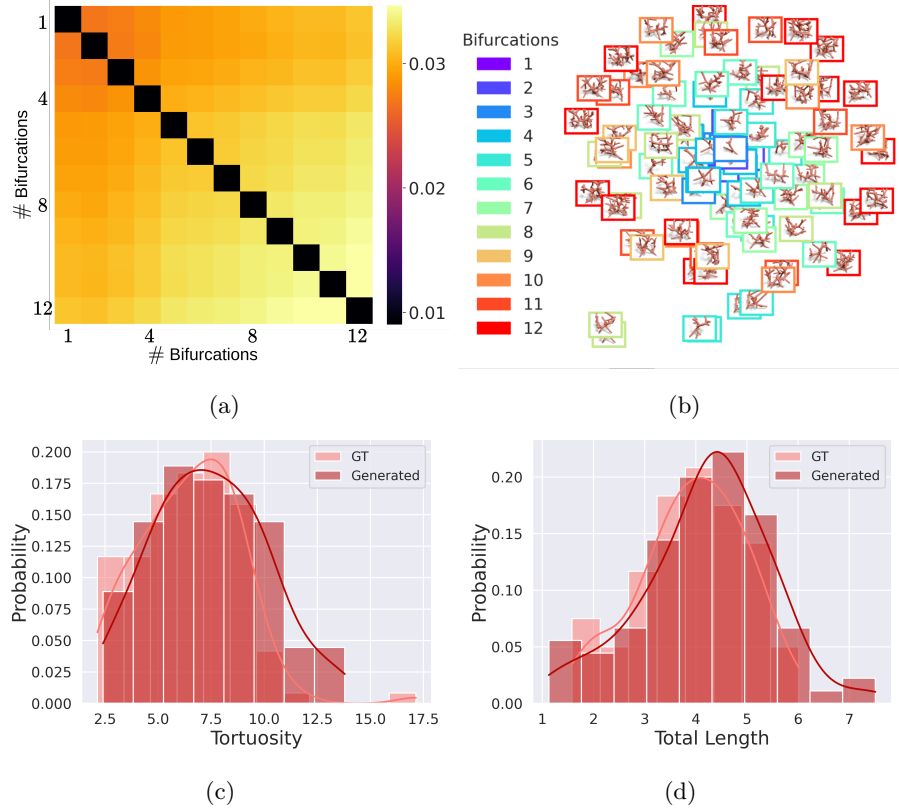


Fig. 6: **Tree Statistics** computed on VasuSynth. (a) Distance between INRs of trees with similar bifurcations is low and increases as the complexity increases; (b) t-SNE plot of the space of trees as INRs; (c-d) Histograms of tortuosity, total length and average radius for ground truth and generated trees.

Table 2: Quantitative results on tree structures present in different medical imaging modalities represented using INRs. We report the relative percentage error (%) between the reconstructed signal and ground truth.

Modality	Rel. Error (%)	Input Size (MB)	INR Size (MB)
DRIVE (RGB) [32]	0.018	0.37 \pm 0.0055	0.066 \downarrow \times 5.60
DRIVE (Mask) [32]	1.204	0.02 \pm 0.0013	0.003 \downarrow \times 6.60
BraTS [20]	0.039	68.11 \pm 0.00	0.753 \downarrow \times 90.45
HAN-Seg [26]	5.627	12.1 \pm 1.55	0.630 \downarrow \times 19.20

Table 3: Quantitative results for novel tree generation on VascuSynth. \uparrow : higher is better; \downarrow : lower is better. For 1-NNA, 50% is ideal.

Metric	Value
MMD \downarrow	13.36 \pm 8.37
COV \uparrow	0.46 \pm 0.11
1-NNA (%) \downarrow	87.49 \pm 8.99

similar INRs. As the difference in branch count increases, the distance increases, visible by the color gradient in Fig. 6a. Furthermore, for a constant difference in branch count (e.g., $|3 - 4| = 1 = |8 - 9|$), the dissimilarity is more pronounced with larger branch counts, despite the numerical difference being the same.

Tree Synthesis. Training a DDM on INR-based trees results in a model that captures the tree distribution and can be sampled using DDIM [31] to generate novel trees. First, we show qualitative results of novel INR-based trees generated by our model in Figure 8. Next, since the evaluation of unconditional generation of tree-like structures can be challenging due to lack of direct correspondence to ground truth data, we follow [4,5] to quantitatively assess the samples generated using diffusion. Specifically, we use minimum matching distance (MMD), coverage (COV), and 1-nearest neighbor accuracy (1-NNA) to measure quality, diversity, and plausibility, respectively. Moreover, since the generations are random, we run the evaluation 3 times and report the mean values in Table 3. Additionally, we follow [6,37] and report vessel-based metrics in Figs. 6c and 6d, where we see high similarities in tortuosity and centerline length histograms between ground truth training data and synthesized trees.

INR-based Image Segmentation. We present two proof-of-concept experiments on how INR representation is leveraged to perform vessel tree segmentation. Figure 7 depicts the evolution of INR-represented segmentation masks as they gradually fit to target vasculature in both CoW and WBMRA.

Implementation Details. We used PyTorch and ADAM optimizer with learning rate $\alpha = 10^{-3}$ for all experiments. To fit INRs to trees, we optimize the MLP architecture in [21] for max 5k iterations on 12GB NVIDIA GeForce GTX TITAN X GPU, taking \sim 1 minute/tree. We train the diffusion model (128-D hidden size, 4 layers, and 4 self-attention modules, as in [24]) on 48GB NVIDIA A40 GPU for 6k epochs, with batch size 8, $0.9 \times \alpha$ decay per 200 epoch, 1k diffusion time-steps, linear noise scheduler $\in [10^{-4}, 10^{-2}]$, taking \approx 3 days.

4 Conclusion

We presented the first work to use implicit neural fields for faithful representation of topologically-complex anatomical trees and learnt their distribution




Fig. 7: **Segmentation.** The evolution of INR towards segmenting two vessel trees. Analogous to Mumford-Shah based segmentation [3], we use a piecewise-constant version of the INR (via simple thresholding) as it is being optimized to faithfully represent the original image.




Fig. 8: **Tree Synthesis.** Some novel INRs sampled from the diffusion model, visualized as meshes.

via training a diffusion model in the space of neural fields. We demonstrated quantitatively and qualitatively the advantages of our method: versatility (*e.g.*, 2D/3D; vascular/airway; simple/complex topology), lower memory footprint while achieving highly accurate reconstructions at arbitrary high resolutions; synthesis of plausible trees; and application to segmentation of medical images. Further, our representation is amenable to integration into deep learning pipelines and can be easily transformed into other shape representations. Future work may include integrating our method into more advanced deep segmentation pipelines while encoding semantic annotations of tree parts.

5 Acknowledgements

The authors are grateful to Paula Feldman for the provision of evaluation scripts used for computing tortuosity and other tree-specific metrics. This research was enabled in part by support provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), and the computational resources provided by Simon Fraser University and NVIDIA Corporation.

References

1. Alblas, D., et al.: Going off-grid: Continuous implicit neural representations for 3D vascular modeling. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 79–90. Springer (2022) [1](#), [2](#), [3](#)
2. van Bemmelen, C.M., et al.: Level-set based carotid artery segmentation for stenosis grading. In: MICCAI. pp. 36–43. Springer (2002) [2](#)
3. Chan, T.F., et al.: Active contours without edges. IEEE Transactions on image processing **10**(2), 266–277 (2001) [10](#)
4. Chou, G., et al.: Diffusion-SDF: Conditional generative modeling of signed distance functions. In: CVPR. pp. 2262–2272 (2023) [3](#), [9](#)
5. Erkoç, Z., et al.: HyperDiffusion: Generating implicit neural fields with weight-space diffusion. arXiv preprint arXiv:2303.17015 (2023) [3](#), [4](#), [9](#)
6. Feldman, P., et al.: VesselVAE: Recursive variational autoencoders for 3D blood vessel synthesis. In: MICCAI. pp. 67–76. Springer (2023) [3](#), [9](#)
7. Feragen, A., et al.: Toward a theory of statistical tree-shape analysis. IEEE TPAMI **35**(8), 2008–2021 (2012) [3](#)
8. Galarreta-Valverde, M.A., et al.: Three-dimensional synthetic blood vessel generation using stochastic L-systems. In: Medical Imaging 2013: Image Processing. vol. 8669, pp. 414–419. SPIE (2013) [3](#)
9. Hamarneh, G., et al.: VasuSynth: Simulating vascular trees for generating volumetric image data with ground-truth segmentation and tree analysis. Computerized medical imaging and graphics **34**(8), 605–616 (2010) [3](#), [5](#)
10. Ho, J., et al.: Denoising diffusion probabilistic models. NeurIPS **33**, 6840–6851 (2020) [3](#)
11. Hong, Q., et al.: High-quality vascular modeling and modification with implicit extrusion surfaces for blood flow computations. Computer Methods and Programs in Biomedicine **196**, 105598 (2020) [2](#)
12. Hu, S.M., et al.: Subdivision-based mesh convolution networks. ACM (TOG) **41**(3), 1–16 (2022) [2](#)
13. Kretschmer, J., et al.: Interactive patient-specific vascular modeling with sweep surfaces. IEEE TVCG **19**(12), 2828–2837 (2013) [2](#)
14. Li, H., et al.: Vessels as 4-D curves: Global minimal 4-D paths to extract 3-D tubular surfaces and centerlines. IEEE TMI **26**(9), 1213–1223 (2007) [1](#)
15. Lindenmayer, A., et al.: Mathematical models for cellular interactions in development I. filaments with one-sided inputs. Journal of theoretical biology **18**(3), 280–299 (1968) [1](#)
16. Lo, P., et al.: Extraction of airways from CT (EXACT’09). IEEE TMI **31**(11), 2093–2107 (2012) [5](#)
17. Lorensen, W.E., et al.: Marching cubes: A high resolution 3d surface construction algorithm. In: Seminal graphics: pioneering efforts that shaped the field, pp. 347–353 (1998) [3](#)
18. Lorigo, L.M., et al.: CURVES: Curve evolution for vessel segmentation. MedIA **5**(3), 195–206 (2001) [2](#)
19. Maturana, D., et al.: VoxNet: A 3D convolutional neural network for real-time object recognition. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 922–928. IEEE (2015) [2](#)
20. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). IEEE TMI **34**(10), 1993–2024 (2014) [5](#), [9](#)

21. Mescheder, L., et al.: Occupancy networks: Learning 3D reconstruction in function space. In: CVPR. pp. 4460–4470 (2019) [2](#), [3](#), [4](#), [5](#), [9](#)
22. Niemeyer, M., et al.: Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In: CVPR. pp. 3504–3515 (2020) [3](#)
23. Park, J.J., et al.: DeepSDF: Learning continuous signed distance functions for shape representation. In: CVPR. pp. 165–174 (2019) [2](#), [3](#)
24. Peebles, W., et al.: Learning to learn with generative models of neural network checkpoints. arXiv preprint arXiv:2209.12892 (2022) [4](#), [9](#)
25. Pizaine, G., et al.: Implicit medial representation for vessel segmentation. In: Medical Imaging 2011: Image Processing. vol. 7962, pp. 1184–1190. SPIE (2011) [1](#)
26. Podobnik, G., et al.: HaN-Seg: The head and neck organ-at-risk CT and MR segmentation dataset. Medical physics **50**(3), 1917–1927 (2023) [5](#), [9](#)
27. Prusinkiewicz, P., et al.: Synthetic topiary. In: Annual Conference Series. ACM SIGGRAPH, Addison Wesley. vol. 10 (1994) [2](#)
28. Qi, C.R., et al.: Volumetric and multi-view CNNs for object classification on 3D data. In: CVPR. pp. 5648–5656 (2016) [2](#)
29. Sitzmann, V., et al.: Implicit neural representations with periodic activation functions. NeurIPS **33**, 7462–7473 (2020) [2](#)
30. Sobocinski, J., et al.: The benefits of EVAR planning using a 3D workstation. European Journal of Vascular and Endovascular Surgery **46**(4), 418–423 (2013) [1](#)
31. Song, J., et al.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020) [5](#), [9](#)
32. Staal, J., et al.: Ridge-based vessel segmentation in color images of the retina. IEEE TMI **23**(4), 501–509 (2004) [5](#), [9](#)
33. Talou, G.D.M., et al.: Adaptive constrained constructive optimisation for complex vascularisation processes. Scientific Reports **11**(1), 6180 (2021) [3](#)
34. Tran, K., et al.: Patient-specific computational flow modelling for assessing hemodynamic changes following fenestrated endovascular aneurysm repair. JVS-vascular science **2**, 53–69 (2021) [1](#)
35. Vaswani, A., et al.: Attention is all you need. NeurIPS **30** (2017) [4](#)
36. WHO, et al.: Global health estimates: Leading causes of death (2019) [1](#)
37. Wolterink, J.M., et al.: Blood vessel geometry synthesis using generative adversarial networks. arXiv preprint arXiv:1804.04381 (2018) [3](#), [9](#)
38. Yang, X., et al.: Intra: 3D intracranial aneurysm dataset for deep learning. In: CVPR (2020) [1](#), [5](#)
39. Zamir, M., et al.: Arterial branching within the confines of fractal L-system formalism. The Journal of general physiology **118**(3), 267–276 (2001) [3](#)
40. Zhang, B., et al.: 3DShape2VecSet: A 3D shape representation for neural fields and generative diffusion models. arXiv preprint arXiv:2301.11445 (2023) [3](#)
41. Zhao, M., et al.: Leveraging tree statistics for extracting anatomical trees from 3D medical images. In: CRV. pp. 131–138. IEEE (2017) [3](#)
42. Zhao, M., et al.: Tree-LSTM: using LSTM to encode memory in anatomical tree prediction from 3D images. In: MICCAI Workshop (MLMI). pp. 637–645. Springer (2019) [1](#)