



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Self-supervised 3D Skeleton Completion for Vascular Structures

Jiaxiang Ren¹, Zhenghong Li¹, Wensheng Cheng¹, Zhilin Zou¹, Kicheon Park², Yingtian Pan², and Haibin Ling^{1*}

¹ Department of Computer Science, Stony Brook University
hling@cs.stonybrook.edu

² Department of Biomedical Engineering, Stony Brook University

Abstract. 3D skeleton is critical for analyzing vascular structures with many applications, it is however often limited by the broken skeletons due to image degradation. Existing methods usually correct such skeleton breaks via handcrafted connecting rules or rely on nontrivial manual annotation, which is susceptible to outliers or costly especially for 3D data. In this paper, we propose a self-supervised approach for vasculature reconnection. Specifically, we generate synthetic breaks from confident skeletons and use them to guide the learning of a 3D UNet-like skeleton completion network. To address serious imbalance among different types of skeleton breaks, we introduce three skeleton transformations that largely alleviate such imbalance in synthesized break samples. This allows our model to effectively handle challenging breaks such as bifurcations and tiny fragments. Additionally, to encourage the connectivity outcomes, we design a novel differentiable connectivity loss for further improvement. Experiments on a public medical segmentation benchmark and a 3D optical coherence Doppler tomography (ODT) dataset show the effectiveness of our method. The code is publicly available at <https://github.com/reckdk/SkelCompletion-3D>.

Keywords: 3D skeleton completion · Self-supervised learning.

1 Introduction

Skeletonization is a fundamental step for analyzing curvilinear structures, such as blood vessels and neurons [10, 24, 25, 28]. Based on the topological structure of the skeletons, researchers and doctors can quantitatively assess vascular distribution, density, and morphological changes, which provides valuable insights into diseases, drug effects, and neuroscience studies. For example, the precise reconstruction of vasculature is crucial in the analysis of tumor angiogenesis, reflecting the cancer progression and response [27]. Besides, the rupture of brain vessels is one of the leading causes of stroke, which results in numerous fatalities each year. Identifying the process of vascular remodeling, such as the connections between arteries, has shown great promise in improving the prognosis accuracy and patient outcomes [2, 15].




Fig. 1: Examples of 3D skeleton completion. From left to right: 3D ODT volume, skeleton input, and completed skeleton. Our work focuses on reconnecting the broken skeletons indicated by green arrows.

Despite previous efforts on automatic skeletonization [14, 20, 21, 31], broken skeletons often occur due to image noise and disease-related appearance changes [30], introducing difficulties in quantitative analysis. In the widely-used pipeline from image to segmentation to skeleton, the accuracy of skeletonization relies on the quality of segmentation. In scenarios where under-segmentation error occurs, particularly in the noise-affected areas, it results in incomplete foreground masks and, consequently, leads to fragmented skeletons.

Previous methods for correcting breaks can be roughly categorized as rule-based and learning-based. Rule-based methods [4, 17] deal with broken skeletons directly and usually act as post-processing that detects breaks and then links the fragments. The detection contains a set of handcrafted policies, such as constraints on distance, curvature [4], and curve-fitting [17, 30], to determine whether the two fragments should be connected. Rule-based methods requires minimal, if any, training data requirements. However, the manually defined constraints may not generalize to various complex scenarios. On the other hand, the learning-based methods restore breaks in the preceding segmentation stage to indirectly reconnect skeleton fragments. Various approaches based on graphs [6, 16], topology [11, 26], and loss functions [19, 29] have been proposed to enhance segmentation results. Learning-based methods can acquire more representative features during training and are thus more flexible in dealing with breaks of various types of morphology. However, these methods are based on supervised learning and thus not feasible for the datasets without ground truth annotations.

Recently, self-supervised learning (SSL) [9, 32] addresses the scarcity of ground truth data through a set of training schemes, garnering increasing attention [7]. Integrating the confident and learnable information plays a pivotal role in SSL. Meanwhile, we have observed that, though underperforming in noisy regions, current skeletonization algorithms produce reliable topological structures in most normal regions. Leveraging such reliable structures can benefit the training process to capture more connectivity related features for reconstruction.

Thus inspired, we propose a self-supervised 3D skeleton completion approach to deal with the aforementioned challenges: (1) Our model is trained in a self-supervised manner, eliminating the need for manual annotations. We achieve this by synthesizing breaks from confident skeletons and use them to guide the learning of a 3D UNet-like [5] skeleton completion network. Subsequently, our model utilizes the prior learned skeleton structures for real break completion in the same samples, demonstrating self-adaptiveness. (2) To handle the significant




Fig. 2: (a) Vascular forests. (b) Relationship between skeleton transformations and break types. Solid line: direct synthesis; dashed line: indirect synthesis.

imbalance across various types of skeleton breaks, we propose three skeleton transformations that largely mitigate the long tail issues in synthetic break samples. This improves the model’s completion ability, especially in challenging cases such as bifurcations and tiny fragments. (3) To enhance connectivity outcomes, we design a novel loss function to further improve the performance, which is made differentiable for end-to-end training. Besides, different from previous methods, we decouple skeleton completion from image modalities or segmentation, which relieves the difficulty in learning dissimilar data distributions and filters out intensity noises. Hence our model is capable of reconnecting the broken skeleton from various segmentation approaches, *i.e.*, gradient-based filters, deep learning models, and imperfect manual annotations, for different image modalities, including CT, ultrasound, and ODT.

Comprehensive evaluations are conducted on the popular Medical Segmentation Decathlon (MSD) dataset [3] and a 3D optical coherence Doppler tomography (ODT) dataset. The results show clearly the effectiveness of the proposed method in comparison with previous arts.

2 Methods

2.1 Problem Formulation

The input of our skeleton completion problem contains a 3D volume and initially extracted skeletons organized in a forest $\mathcal{F} := \{\mathcal{T}^1, \mathcal{T}^2, \dots\}$ of trees; each tree $\mathcal{T}^i := \{\mathcal{S}_1^i, \mathcal{S}_2^i, \dots\}$ has multiple skeletons; each skeleton $\mathcal{S}_j^i := [v_{j,1}^i, v_{j,2}^i, \dots]$ has an ordered list of voxel-based nodes; and each node is defined by a 6-tuple

$$v_{j,k}^i := (z_{j,k}^i, y_{j,k}^i, x_{j,k}^i, d_{j,k}^i, sid_{j,k}^i = j, tid_{j,k}^i = i), \quad (1)$$

containing the 3D coordinates, degree, skeleton index, and tree index respectively. The degree is the number of nearby connected voxels. We utilize the 6-tuple for selecting input patches and synthesizing breaks.

The left sub-figure of Fig. 2a illustrates a vascular forest with two trees (\mathcal{T}^1 and \mathcal{T}^2). \mathcal{T}^1 has three skeletons (\mathcal{S}_1^1 , \mathcal{S}_2^1 , and \mathcal{S}_3^1) while \mathcal{T}^2 has only one skeleton \mathcal{S}_1^2 . However, \mathcal{T}^1 will split into three trees (\mathcal{T}^1 , \mathcal{T}^3 , and \mathcal{T}^4) when there is a Bifur-break, as illustrated in the right sub-figure of Fig. 2a. Without loss of




Fig. 3: Training pipeline of the proposed 3D skeleton completion approach. Red voxels are removed for training.

generality, we assume there is a break between \mathbf{v}_a and \mathbf{v}_b (indices are replaced with a and b for simplicity), if the following conditions are met:

1. $\|(z_a, y_a, x_a) - (z_b, y_b, x_b)\| \leq \tau$, for a predefined distance threshold τ .
2. The curvilinear structures near \mathbf{v}_a and \mathbf{v}_b show similar geometric properties, such as direction and curvature.
3. $sid_a \neq sid_b$ and $tid_a \neq tid_b$, *i.e.*, no circle in the tree.

Our task is to reconnect such breaks in the vascular forest.

2.2 Skeleton Transformations

A primary challenge in vascular skeleton completion is the diversity and uncertainty inherent in real-world breaks. Based on location, branch number, and length, we categorize breaks into four types: (1) Mid-breaks (split a skeleton in the middle to two long branches), (2) Tip-break (breaks near the end of a skeleton), (3) Multi-breaks (breaks with more than two branches in a line), and (4) Bifur-breaks (bifurcation breaks with more than two interconnected branches), as shown in Fig. 2b. The naive SSL, *i.e.*, simply removing mid voxels in a skeleton for training, has two critical issues: (1) it can only generate Mid-break and hence has poor generalizability for rare and complex breaks; and (2) the synthesized breaks are always located at the patch center to avoid crossing boundaries, leading to severely performance drop as breaks deviate from the center. To address the aforementioned issues of naive SSL, we propose three types of skeleton transformation as follows.

- **RoI Shifting** (RoI Shift) shifts the patch several voxels away from the original center, improving the robustness against location bias.
- **Skeleton Slipping** (Skel Slip) removes voxels towards the skeleton end rather than the center, enriching the training samples significantly.
- **Skeleton Shattering** (Skel Shatter) randomly removes of discontinuous voxels to simulate severely broken skeletons with multiple fragments.

Fig. 2b illustrates the relationship between skeleton transformations and break types. These transformations broaden the coverage of synthesized breaks, effectively mitigating long tail issues. Furthermore, both RoI Shift and Skel Slip relax the model’s attention away from the exact center, restoring breaks anywhere in the patch. Consequently, our model waives the need for explicit break detection. Notably, our model can autonomously reconnect most breaks in a volume via

inferring at all endpoints. This fully automatic approach is efficient as locating endpoints proves more practical than detecting breaks.

2.3 Self-supervised Learning and Connectivity Loss

Given a dataset with reliable skeletons in most regions except the skeleton breaks, we introduce a self-supervised skeleton completion approach to reconnect these breaks. Fig. 3 illustrates the training pipeline. Initially, we randomly sample 3D patches $\mathbf{X} \in \{0, 1\}^{W \times W \times W}$ with long skeletons. Next, we apply the proposed skeleton transformations on \mathbf{X} to generate input patches $\{\mathbf{X}'_1, \mathbf{X}'_2, \dots\}$ with corresponding binary masks $\{\mathbf{M}_1, \mathbf{M}_2, \dots\}$ (where 1 indicates removal for training). Lastly, we use these synthetic break samples to train the 3D skeleton completion network f , which is derived from the 3D UNet [5], as shown in Fig. 3. The model output $\mathbf{Y} = f(\mathbf{X}') \in [0, 1]^{W \times W \times W}$ is a confidence map of missing skeletons. A higher response means the corresponding voxel is more likely to be a skeleton.

General segmentation losses such as *DiceLoss* and *TverskyLoss* [22] focus on the accuracy of all voxels while overlooking the connectivity of prediction. In this paper, we propose a customized ℓ_1 loss that penalizes breaks on skeleton and thus encourages connectivity:

$$\begin{aligned} \mathcal{L}_{\text{conn}} &= \mathcal{L}_{\text{R}} + \omega \mathcal{L}_{\text{P}}, \text{ where } \mathcal{L}_{\text{R}} = \mathbf{M}(\|\mathbf{X}' - \mathbf{Y}\|_1 + \lambda * \delta(\mathbf{Y} - 0.5)) / \|\mathbf{M}\|, \\ &\mathcal{L}_{\text{P}} = (\mathbf{1} - \mathbf{M})\|\mathbf{X}' - \mathbf{Y}\|_1 / \|\mathbf{1} - \mathbf{M}\|, \end{aligned} \quad (2)$$

where $\delta(\cdot)$ is the indicator function, ω and λ are weights of losses. The differentiable *ConnLoss* is the weighted sum of the restoring loss \mathcal{L}_{R} and the preservation loss \mathcal{L}_{P} . \mathcal{L}_{R} comprises two items: (1) the regular ℓ_1 loss to guide the prediction towards the target; and (2) an extra penalty for false negative voxels. \mathcal{L}_{P} maintains the integrity of non-break area while mitigating false positives. Due to the extreme class imbalance, both \mathcal{L}_{R} and \mathcal{L}_{P} are normalized by their respective sizes. Once training is complete, our model runs without additional user intervention, thereby reducing uncertainty. During inference, we binarize the output \mathbf{Y} and apply the skeletonization algorithm [14] to obtain the final skeleton.

3 Experiments and Results

MSD Dataset. MSD [3] contains 303 CT volumes (242 for training and 61 for testing) with mask annotations for Vessel segmentation. In the test set, we manually annotate all patches containing at least two endpoints. The annotation of each patch is a quintuple $(c, \mathbf{v}_a, \mathbf{v}_b, \mathbf{v}_a^*, \mathbf{v}_b^*)$, where c is the binary break label, $(\mathbf{v}_a, \mathbf{v}_b)$ and $(\mathbf{v}_a^*, \mathbf{v}_b^*)$ denote respectively the voxels of the two nearest endpoints and manually verified endpoints. For the trade-off between annotation confidence and the maximum break length, we set $\tau = 10$ that works for most cases. After the experts' verification, there are 152 break patches and 143 non-break ones.

ODT Dataset. We collect 13 ODT volumes for 3D brain vasculature. The dataset presents greater challenges due to severe noise and densely distributed

capillaries. Consequently, there are more incomplete vessels while vascular annotation is prohibitively expensive. We observed that, despite heavy noise, the Frangi filter [8] and skeletonization algorithm [14] produce satisfactory results for most large and clear vessels, which are utilized for training. Additionally, we annotate 600 patches randomly selected from four volumes using a similar procedure as MSD. Specifically, there are 154 Mid-breaks, 20 Tip-breaks, 13 Multi-breaks, 24 Bifur-breaks, and 389 patches without break. Additional details regarding the dataset and annotation can be found in Suppl.

Implementation and Baselines. Our method is implemented on PyTorch and trained using an NVIDIA 3090 GPU. All models are trained with the Adam [13] optimizer for 200 epochs with a batch size of 240. The loss weights ω and λ are set to 1.0 and 2.0 respectively. The initial learning rate is 10^{-3} and halves if no improvement observed for 20 epochs. The input patch size is set to $16 \times 16 \times 16$ for MSD and $32 \times 32 \times 32$ for ODT, based on their statistics of breaks. All skeletons are acquired with the skeletonization method [14]. We use *napari* [1] for visualization and *Skan* [18] for skeleton analysis. The distance threshold τ is 10 for both datasets, as suggested by experts. We use the 3D Frangi filter with scales ranging from 1 to 5 and a scale step of 0.5 to extract vessel masks. The hyperparameters in the Frangi filter are set as $\alpha = 0.5$, $\beta = 0.5$, and $\gamma = 15.0$.

The "Distance-based" baseline simply detects breaks based on the distance between two endpoints, and connects detected points with the shortest path. Vess. Comp. [4] is the state-of-the-art skeleton completion approach that detects breaks using both distance and curvature similarity. nnUNet [12] is a state-of-the-art medical image segmentation model, and we convert its prediction into skeletons for comparison. Hyperparameters are fine-tuned using grid search to achieve optimal performance. nnUNet results on ODT are not reported due to the lack of manual annotations for fine-tuning.

3.1 Quantitative Experiment

We define the following metrics for skeleton completion:

$$\begin{aligned} \text{TP} &= \sum_i c_i * \kappa(\mathbf{v}_{a(i)}^*, \mathbf{v}_{b(i)}^*), & \text{FN} &= \sum_i c_i * (1 - \kappa(\mathbf{v}_{a(i)}^*, \mathbf{v}_{b(i)}^*)), \\ \text{TN} &= \sum_i (1 - c_i) * (1 - \kappa(\mathbf{v}_{a(i)}, \mathbf{v}_{b(i)})), & \text{FP} &= \sum_i (1 - c_i) * \kappa(\mathbf{v}_{a(i)}, \mathbf{v}_{b(i)}), \end{aligned} \quad (3)$$

where i is the patch index, c_i is its class label, and $\kappa(\mathbf{v}_a, \mathbf{v}_b)$ the connectivity function that returns 1 if \mathbf{v}_a and \mathbf{v}_b are connected and 0 otherwise. For a break patch, the predicted skeleton is true positive if it connects the two manually-verified endpoints \mathbf{v}_a^* and \mathbf{v}_b^* ; otherwise, it is false negative. For a patch without break, the predicted skeleton is true negative if it does not connect the two nearest endpoints \mathbf{v}_a and \mathbf{v}_b ; otherwise, it is false positive. Recall, precision, accuracy, and F1 score are measured accordingly. Due to the significant imbalance in break types, F1 score is preferred over other metrics. The point-to-point metrics [23] requiring costly ground truth are beyond the scope of our work.

Table 1: Performance (%) on the **MSD** and **ODT** datasets. Rec: recall; Prec: precision; Acc: accuracy. nnUNet* is skeletonized prediction of pretrained nnUNet. nnUNet*+Ours is our SSL model trained with the prediction of nnUNet.

	MSD				ODT			
	Rec	Prec	Acc	F1	Rec	Prec	Acc	F1
Distance-based	84.9	47.4	43.7	60.9	86.2	63.8	80.6	73.4
Vess. Comp. [4]	61.2	49.0	47.1	54.4	83.8	72.9	75.6	78.0
nnUNet* [12]	32.4	93.6	64.3	48.1	-			
nnUNet*+Ours	93.4	76.8	82.0	84.3	-			
Ours	92.1	80.0	84.0	85.6	93.7	79.5	90.6	86.0

Table 2: Ablation study of the proposed skeleton transformations on the **MSD** dataset.

RoI Shift	Skel Slip	Skel Shatter	Accuracy	F1
			58.5	68.7
✓			82.0	83.2
	✓		83.3	85.2
✓	✓		83.7	85.3
✓	✓	✓	84.0	85.6

Table 3: Ablation study of losses on the **MSD** dataset.

Loss Function	F1	
Skel Aug.	✓	
<i>Tversky</i> _{$\alpha=0.001$}	64.6	80.3
Normalized MSE	67.5	82.3
Normalized ℓ_1	67.5	82.6
<i>ConnLoss</i> (Ours)	72.9	85.6

Tab. 1 reports the main results, and more details can be found in Suppl. On **MSD**, our method significantly outperforms the other methods with accuracy of 84.0% and F1 score of 85.6%. Note that nnUNet is not designed for break completion and encounters challenges in handling significant noise, resulting in relatively low recall. We also use imperfect skeletons from nnUNet to train our self-supervised model, achieving a significant 17.7% improvement in accuracy, thus validating the effectiveness of our method. On **ODT**, the morphological properties are not guaranteed due to severe noises. As a result, the Distance-based method and Vess. Comp. do not perform well. Additional examples illustrating the poor performance of these baseline methods can be found in Suppl. Our method surpasses the state-of-the-art by 8.0% in F1 score. The challenging Multi-breaks and Bifur-breaks bring two issues arise to rule-based methods: (1) it is nontrivial to dynamically adjust the distance threshold, and (2) the loop avoidance demands additional information such as skeleton indices. In contrast, our method implicitly learns the skeleton structures without extra input, achieving 92.3% and 70.8% precision on Multi-breaks and Bifur-breaks respectively.

Ablation Study of Skeleton Transformations. Tab. 2 shows the ablation study of the proposed skeleton transformations. The vanilla self-supervised model has relatively low accuracy; both RoI Shift and Skel Slip significantly boost the performance; and Skel Shatter achieves the best performance.

Ablation Study of Loss Functions. Tab. 3 reports the study of different losses. Regular losses, including *DiceLoss*, MSE, and ℓ_1 fail during the training under the extremely imbalanced class ratio. By contrast, our *ConnLoss* outper-




Fig. 4: Skeleton completion on MSD. (a) The volume with breaks, (b) Bifur-break, (c) Multi-break and our results (in red). Bounding boxes indicate input.




Fig. 5: Hard case results. From left to right: input, Distance-based, Vess. Comp., ours, and reference ODT. Green: break; red: false positive; blue: false negative.

forms other losses, highlighting its superiority. Furthermore, the skeleton transformations boost the performance for all losses. More analysis are in the Suppl.

3.2 Qualitative Experiment

Completion of Various Breaks on MSD. Fig. 4 illustrates examples of skeleton completion on the MSD dataset. In Fig. 4a, disconnected skeletons are encoded in different colors. There are five breaks and one non-break, marked with red and yellow points respectively. Our model completes all five breaks without significantly changing original structures. Skeleton reconstruction at bifurcations is still a challenging task due to unpredictable branch numbers and morphology. Fig. 4b illustrates that our model restores the bifurcation. Fig. 4c illustrates that our model repairs a ruptured skeleton with multiple tiny fragments, indicating that our model is able to discern such structures ignored by previous methods.

Challenging Cases. Fig. 5 illustrates four hard cases. The yellow, orange, and cyan voxels are the prediction of Distance-based, Vess. Comp., and ours. Fig. 5a contains two skeletons from separate vessels so it is not a break. However, both Distance-based and Vess. Comp. produce false positives due to the close proximity and similar shapes of the endpoints. In contrast, our method correctly rejects it. In Fig. 5b, Distance-based wrongly connects the two nearest endpoints and misses the real break, but our method precisely completes the break. In Fig. 5c, only our method reconnects the large break. In Fig. 5d, Distance-based incorrectly links the two nearest endpoints while ignoring the two breaks located closely. Vess. Comp. rejects both breaks due to inconsistent curvatures. Our method restores both breaks, showing the capability for handling complex cases.

4 Conclusion

We propose a self-supervised solution for 3D vasculature reconnection. Our solution is highlighted by three novel skeleton transformations to address the significant imbalance among different types of skeleton breaks during the sample synthesis. Moreover, a novel connectivity loss is designed for further improvement. Experiments on two real-world datasets show clearly the effectiveness of our method.

Acknowledgments. This work was partially supported by NIH grants 1R21DA057699 (HL, YP), 1RF1DA048808 (YP) and 2R01DA029718 (YP), and partially supported by NSF grants 2006665 (HL) and 2128350 (HL).

Disclosure of Interests. The authors declare no competing interests.

References

- napari contributors (2019): napari: a multi-dimensional image viewer for python (Nov 2022), <https://doi:10.5281/zenodo.3555620> **6**
- Abd Aziz, A., Izhar, L.I., Asirvadam, V.S., Tang, T.B., Ajam, A., Omar, Z., Muda, S.: Detection of collaterals from cone-beam ct images in stroke. *Sensors* **21**(23), 8099 (2021) **1**
- Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., et al.: The medical segmentation decathlon. *Nature Communications* **13**(1), 1–13 (2022) **3, 5**
- Chen, C., Zhou, K., Guo, X., Wang, Z., Xiao, R., Wang, G.: Cerebrovascular segmentation in phase-contrast magnetic resonance angiography by multi-feature fusion and vessel completion. *Computerized Medical Imaging and Graphics* **98**, 102070 (2022) **2, 6, 7**
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *MICCAI*. pp. 424–432. Springer (2016) **2, 5**
- Dmitriev, K., Parag, T., Matejek, B., Kaufman12, A.E., Pfister, H.: Efficient correction for em connectomics with skeletal representation. In: *BMVC* (2018) **2**
- Ericsson, L., Gouk, H., Loy, C.C., Hospedales, T.M.: Self-supervised representation learning: Introduction, advances, and challenges. *IEEE Signal Processing Magazine* **39**(3), 42–62 (2022) **2**
- Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: *MICCAI* (1998) **6**
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *CVPR*. pp. 16000–16009 (2022) **2**
- Helmstaedter, M., Briggman, K.L., Denk, W.: High-accuracy neurite reconstruction for high-throughput neuroanatomy. *Nature Neuroscience* **14**(8), 1081–1088 (2011) **1**
- Hu, X., Li, F., Samaras, D., Chen, C.: Topology-preserving deep image segmentation. *NeurIPS* **32** (2019) **2**
- Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021) **6, 7**

13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015) [6](#)
14. Lee, T.C., Kashyap, R.L., Chu, C.N.: Building skeleton models via 3-d medial surface axis thinning algorithms. *CVGIP: Graphical Models and Image Processing* **56**(6), 462–478 (1994) [2](#), [5](#), [6](#)
15. Liu, J., Wang, Y., Akamatsu, Y., Lee, C.C., Stetler, R.A., Lawton, M.T., Yang, G.Y.: Vascular remodeling after ischemic stroke: mechanisms and therapeutic potentials. *Progress in Neurobiology* **115**, 138–156 (2014) [1](#)
16. Matejek, B., Haehn, D., Zhu, H., Wei, D., Parag, T., Pfister, H.: Biologically-constrained graphs for global connectomics reconstruction. In: CVPR. pp. 2089–2098 (2019) [2](#)
17. Mou, L., Chen, L., Cheng, J., Gu, Z., Zhao, Y., Liu, J.: Dense dilated network with probability regularized walk for vessel detection. *TMI* **39**(5), 1392–1403 (2019) [2](#)
18. Nunez-Iglesias, J., Blanch, A.J., Looker, O., Dixon, M.W., Tilley, L.: A new python library to analyse skeleton images confirms malaria parasite remodelling of the red blood cell membrane skeleton. *PeerJ* **6**, e4312 (2018) [6](#)
19. Oner, D., Osman, H., Koziński, M., Fua, P.: Enforcing connectivity of 3d linear structures using their 2d projections. In: MICCAI. pp. 591–601 (2022) [2](#)
20. Saeed, K., Tabedzki, M., Rybnik, M., Adamski, M.: K3m: A universal algorithm for image skeletonization and a review of thinning techniques. In: *International Journal of Applied Mathematics and Computer Science* (2010) [2](#)
21. Saha, P.K., Borgfors, G., di Baja, G.S.: A survey on skeletonization algorithms and their applications. *Pattern Recognition Letters* **76**, 3–12 (2016) [2](#)
22. Salehi, S.S.M., Erdogmus, D., Gholipour, A.: Tversky loss function for image segmentation using 3d fully convolutional deep networks. In: MLMI. pp. 379–387 (2017) [5](#)
23. Schaap, M., Metz, C.T., van Walsum, T., van der Giessen, A.G., Weustink, A.C., Mollet, N.R., Bauer, C., Bogunović, H., Castro, C., Deng, X., et al.: Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Medical Image Analysis* **13**(5), 701–714 (2009) [6](#)
24. Schikora, J., Kiwatrowski, N., Förster, N., Selbach, L., Ostendorf, F., Pallapies, F., Hasse, B., Metzdorf, J., Gold, R., Mosig, A., et al.: A propagated skeleton approach to high throughput screening of neurite outgrowth for in vitro parkinson’s disease modelling. *Cells* **10**(4), 931 (2021) [1](#)
25. Selle, D., Preim, B., Schenk, A., Peitgen, H.O.: Analysis of vasculature for liver surgical planning. *TMI* **21**(11), 1344–1357 (2002) [1](#)
26. Shi, T., Ding, X., Zhou, W., Pan, F., Yan, Z., Bai, X., Yang, X.: Affinity feature strengthening for accurate, complete and robust vessel segmentation. *JBHI* (2023) [2](#)
27. Vakoc, B.J., Fukumura, D., Jain, R.K., Bouma, B.E.: Cancer imaging by optical coherence tomography: preclinical progress and clinical potential. *Nature Reviews Cancer* **12**(5), 363–368 (2012) [1](#)
28. You, J., Volkow, N.D., Park, K., Zhang, Q., Clare, K., Du, C., Pan, Y.: Cerebrovascular adaptations to cocaine-induced transient ischemic attacks in the rodent brain. *JCI insight* **2**(5) (2017) [1](#)
29. Yu, W., Zheng, H., Zhang, M., Zhang, H., Sun, J., Yang, J.: Break: Bronchi reconstruction by geodesic transformation and skeleton embedding. In: ISBI (2022) [2](#)
30. Zhang, J., Bekkers, E., Chen, D., Berendschot, T.T., Schouten, J., Pluim, J.P., Shi, Y., Dashtbozorg, B., ter Haar Romeny, B.M.: Reconnection of interrupted

- curvilinear structures via cortically inspired completion for ophthalmologic images. *TBME* **65**(5), 1151–1165 (2018) [2](#)
31. Zhang, T.Y., Suen, C.Y.: A fast parallel algorithm for thinning digital patterns. *Communications of the ACM* **27**(3), 236–239 (1984) [2](#)
 32. Zhou, Z., Sodha, V., Pang, J., Gotway, M.B., Liang, J.: Models genesis. *Medical Image Analysis* **67**, 101840 (2021) [2](#)