Semi-supervised Tubular Structure Segmentation with Cross Geometry and Hausdorff Distance Consistency

Ruiyun Zhu[1], Masahiro Oda[2,1], Yuichiro Hayashi[1], Takayuki Kitasaka[3], and Kensaku Mori[1,2]

[1] Graduate School of Informatics, Nagoya University, Furo-cho, Chikusa-ku, Nagoya, Aichi, Japan
rzhu@mori.m.is.nagoya-u.ac.jp, kensaku@is.nagoya-u.ac.jp
[2] Information Technology Center, Nagoya University, Furo-cho, Chikusa-ku, Nagoya, Aichi, Japan
[3] School of Information Science, Aichi Institute of Technology, 1247 Yachigusa, Yakusa-cho, Toyota, Aichi, Japan

Abstract. This study introduces a novel semi-supervised method for 3D segmentation of tubular structures. Complete and automated segmentation of complex tubular structures in medical imaging remains a challenging task. Traditional supervised deep learning methods often demand a tremendous number of annotated data to train the deep model, with the high cost and difficulty of obtaining annotations. To address this, a semi-supervised approach could be a viable solution. Segmenting complex tubular structures with limited annotated data remains a formidable challenge. Many semi-supervised techniques rely on pseudo-labeling, which involves generating labels for unlabeled images based on predictions from a model trained on labeled data. Besides, several semi-supervised learning methods are proposed based on data-level consistency, which enforces consistent predictions by applying perturbations to input images. However, these methods tend to overlook the geometric shape characteristics of the segmentation targets. In our research, we introduce a task-level consistency learning approach that incorporates cross geometry consistency and the Hausdorff distance consistency, taking advantage of the geometric shape properties of both labeled and unlabeled data. Our deep learning model generates both a segmentation map and a distance transform map. By applying the proposed consistency, we ensure that the geometric shapes in both maps align closely, thereby enhancing the accuracy and performance of tubular structure segmentation. We tested our method on airway segmentation in 3D CT scans, where it outperformed the recent state-of-the-art methods, showing an 88.4% tree length detected rate, 82.8% branch detected rate, and 89.7% precision rate.

Keywords: semi-supervised learning · tubular structure segmentation · cross geometry consistency · Hausdorff distance consistency.

1 Introduction

The precise extraction of tubular structures, such as airways and blood vessels, is crucial for many computer-aided diagnosis systems. However, accurate and reliable segmentation of these structures in medical imaging remains a significant challenge. In the last decade, deep learning techniques have had a profound impact on medical image analysis. While many fully supervised deep learning models have reached state-of-the-art performance in tubular structure segmentation [22, 23], they are heavily reliant on the amount of labeled data. Yet the task of annotating tubular structures in medical images, such as airways in chest CT scans, is notably time-consuming and requires extensive expertise.

To mitigate the issue, many studies on semi-supervised learning (SSL) have been proposed. SSL methods leverage both labeled and unlabeled data to improve segmentation performance. The pseudo-labeling technique is an intuitive and common strategy for generating pseudo-labels for unlabeled samples based on the prediction results from a model trained on labeled data [26, 6]. Besides the pseudo-labeling approach, unlabeled data can be utilized via consistency learning, which encompasses two categories: data-level and task-level consistency learning. Data-level consistency learning promotes the stability of model predictions for input images, even when these images undergo various perturbations. French et al. [5] introduced an input perturbation by applying the CutMix technique [25] to input images. Ouali et al. [15] proposed several feature perturbations to enforce the consistency of an input image and its perturbed versions. Although the abovementioned methods achieved good performance in semantic segmentation through various perturbations, they ignore the geometric information of segmentation objectives.

Task-level consistency learning emphasizes establishing regularization at the task level through the addition of auxiliary tasks, enabling the utilization of geometric information [13, 2]. Luo et al. [12] blended the level set function regression task with the segmentation task, establishing a dual-task consistency framework for SSL. Their method enforces the predicted segmentation maps to be consistent with the predicted signed distance transform maps. Liu et al. [9] proposed a shape-aware multi-task framework that contained segmentation, signed distance map prediction, and organ contour prediction. Liu et al. [10] introduced a dual-view network to predict two pairs of a distance transform (DT) map and a segmentation map. They applied a geometry consistency to the predicted maps.

Previous studies in task-level consistency learning have shown impressive results in segmentation by leveraging geometric information from auxiliary tasks, yet have ignored the potential of geometric information from the segmentation task itself. We propose that geometric information can be derived not only from auxiliary tasks but also directly from the segmentation task. Predicted segmentation maps can be transformed to distance transform maps, offering a deeper understanding of geometric characteristics. Our approach introduces a method to harness geometric information through cross geometry consistency and Hausdorff distance consistency from both segmentation maps and distance transform maps, aiming to improve segmentation performance.

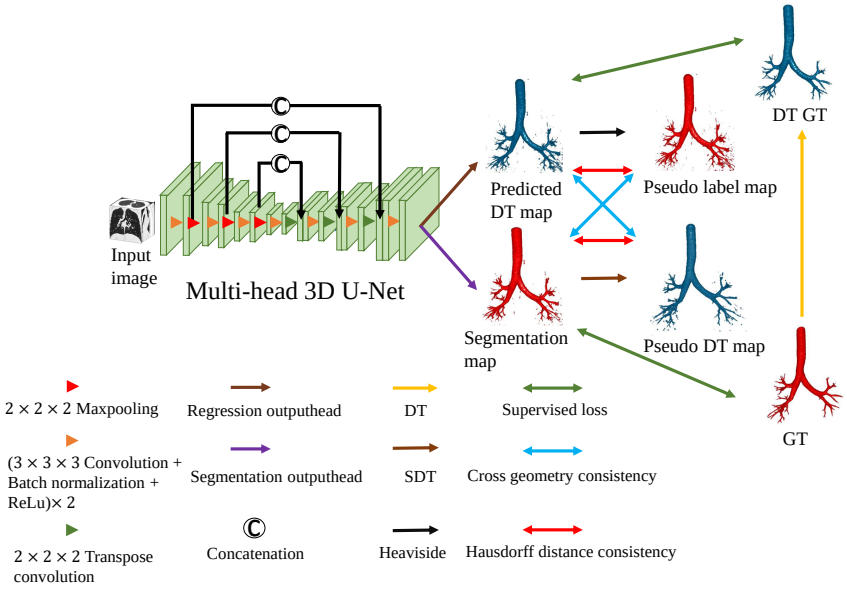

Fig. 1: The SSL framework of our method.

The contributions of this work are summarised as follows: **1)** We design a task-level framework for SSL. **2)** We introduce a cross geometry and a Hausdorff distance consistency learning method which could leverage geometric information from both the segmentation task and the auxiliary task. Our method achieved an 88.4% tree length detected rate and 82.8% branch detected rate, outperforming previous methods.

2 Method

2.1 Overview

This study aims to develop a SSL framework to increase the accuracy of segmenting tubular structures. We emphasize a consistency learning method that utilizes additional information by introducing an auxiliary task. For this purpose, we use the DT regression task, chosen for its ability to reflect the geometric features of the segmentation objectives [14]. Our task-level SSL framework is shown in Fig. 1. Our deep model is a multi-head convolutional neural network, built upon the architecture of the 3D U-Net [3]. Our network has two output heads with respect to the main segmentation task and the auxiliary DT regression task.

2.2 Cross geometry consistency

The concept of cross geometry consistency is straightforward yet effective. A DT map and segmentation map should mutually represent the same geometric shape

of the segmentation objective. This implies that a predicted segmentation map can serve as an unsupervised reference for the predicted DT map and vice versa. Our cross geometry consistency is founded on this mutual referencing, including two kinds of consistencies: the segmentation-reference consistency and the DT-reference consistency. Let $\mathbf{P} = \{\mathbf{p}_i | i = 1, ..., N\}$ and $\mathbf{R} = \{\mathbf{r}_i | i = 1, ..., N\}$ represent the sets of predicted segmentation maps and predicted DT maps, respectively. The segmentation-reference consistency loss is defined as

$$L_{seg-ref}(\mathbf{p}_i, \mathbf{r}_i) = \sum_{i=1}^{|\mathbf{P}|} \|\mathbf{H}(\mathbf{r}_i) - \mathbf{p}_i\|_2, \tag{1}$$

where $|\mathbf{P}|$ represents the number of the total predicted segmentation maps, as same as the predicted DT maps. $\mathbf{H}(\cdot)$ is a smoothed Heaviside function that transforms a predicted DT map to a binarized pseudo-label map in a differentiable way, producing the reference of the predicted segmentation map. We utilize the hyperbolic tangent function as the smoothed Heaviside function, which is written as

$$\mathbf{H}(\mathbf{r}_i) = \tanh(\gamma \mathbf{r}_i), \tag{2}$$

where γ is a large constant to approximate a Heaviside function.

Similarly, the DT-reference consistency loss is defined as

$$L_{DT-ref}(\mathbf{p}_i, \mathbf{r}_i) = \sum_{i=1}^{|\mathbf{P}|} \|\mathbf{r}_i - D(\mathbf{p}_i)\|_2, \tag{3}$$

where $D(\cdot)$ denotes the soft distance transform (SDT) to produce pseudo-DT maps. In contrast to the differentiable soft binarization, the differentiable SDT is more challenging due to the calculation of finding the minimal distance to the boundary of the segmentation objective. Inspired by [17, 27], we present a differentiable DT calculation. Firstly, we binarize the predicted segmentation maps with another smoothed Heaviside function $\hat{\mathbf{H}}(\cdot)$ using a sigmoid function, similarly to $\mathbf{H}(\cdot)$. We calculate a foreground boundary map through the function $B(\mathbf{p}_i) = Dilation(\hat{\mathbf{H}}(\mathbf{p_i})) - \hat{\mathbf{H}}(\mathbf{p_i})$ with morphological operations. $Dilation(\cdot)$ is a soft dilation operation that is simulated using a $3 \times 3 \times 3$ max pooling function [19]. Let $\mathbf{B} = \{b_i | i = 1, .., m\}$ denotes the set of boundary voxels in a boundary map. Given a foreground voxel f in a segmentation map, $d(f, b_1), ..., d(f, b_m)$ represents the distances from the foreground voxel f to boundary voxels. $d(\cdot, \cdot)$ denotes the Euclidean distance. Via the log-sum-exponential [17], the minimal distance function $M(f, b_i)$ is defined as

$$\begin{aligned} &M(f, b_i) = \min d(f, b_1), ..., d(f, b_m) \\ &= \lim_{\lambda \to 0} -\lambda \log \left(exp\left(-\tfrac{d(f,b_1)}{\lambda} + ... + exp\left(-\tfrac{d(f,b_m)}{\lambda} \right) \right) \right), \end{aligned} \tag{4}$$

where λ approaching 0 from the left. In practice, the λ could be a small negative number (e.g., -0.1). The element-wise minimal distance calculation could be implemented with convolutional operations. We utilize a 3D convolutional

kernel $\mathbf{K}$ that each element is initialized with a Euclidean distance value from the element to the centroid of the $\mathbf{K}$. The centroid locate at the average position of all the points along three dimensions. We set the kernel size to $33 \times 33 \times 33$ to ensure every voxel inside the foreground region could touch the boundary. Following [27], a SDT map could be calculated by

$$D(\mathbf{p}_i) = -\lambda \log(B(\mathbf{p}_i) * exp\left(-\frac{\mathbf{K}}{\lambda}\right), \tag{5}$$

where $*$ denotes the convolution operation using the kernel $exp\left(-\frac{\mathbf{K}}{\lambda}\right)$. $exp\left(-\frac{\mathbf{K}}{\lambda}\right)$ is an element-wise operation which applies the exponential function to each element of $\mathbf{K}$. An example of our SDT map is shown in Fig. 2. Our cross geometry consistency comprises the segmentation-reference consistency and the DT-reference consistency to improve segmentation performance, which is written as

$$L_{cross}(\mathbf{P}, \mathbf{R}) = L_{seg-ref}(\mathbf{P}, \mathbf{R}) + L_{DT-ref}(\mathbf{P}, \mathbf{R}). \tag{6}$$

2.3 Hausdorff distance consistency

The Hausdorff distance is calculated between the boundaries of the estimated segmentation and the ground-truth segmentation. It demonstrates the boundary information of the segmentation objectives. Following [7], in image segmentation, the Hausdorff distance metric could be written as

$$L_{hd}(\mathbf{p}_i, \mathbf{g}_i) = \frac{1}{|\mathbf{P}|} \sum_{i=1}^{|\mathbf{P}|} ((\mathbf{p}_i - \mathbf{g}_i)^2 \circ (\mathbf{p}_i^{dt} + \mathbf{g}_i^{dt})), \tag{7}$$

where $\circ$ denotes Hadamard product. $\mathbf{p}_i$ and $\mathbf{g}_i$ represent the $i-$th predicted segmentation map and the corresponding ground truth. $\mathbf{p}_i^{dt}$ and $\mathbf{g}_i^{dt}$ represent the corresponding DT maps of the predicted segmentation map and the ground truth.

In our Hausdorff distance consistency, the pseudo-label maps and pseudo-DT maps are leveraged as the ground truth maps and the DT ground truth maps, respectively. The proposed consistency enables the predicted segmentation maps and DT maps to mutually supervise the boundary segmentation. Our Hausdorff distance consistency is defined as

$$L_{hdc}(\mathbf{p}_i, \mathbf{r}_i) = \frac{1}{|\mathbf{P}|} \sum_{i=1}^{|\mathbf{P}|} ((\mathbf{p}_i - H(\mathbf{r}_i))^2 \circ (\mathbf{r}_i + D(\mathbf{p}_i))). \tag{8}$$

2.4 Semi-supervised training

Our semi-supervised training framework integrates supervised and unsupervised training. The labeled data is utilized for both segmentation and DT regression tasks in a supervised manner, with L_{seg} representing the loss for segmentation

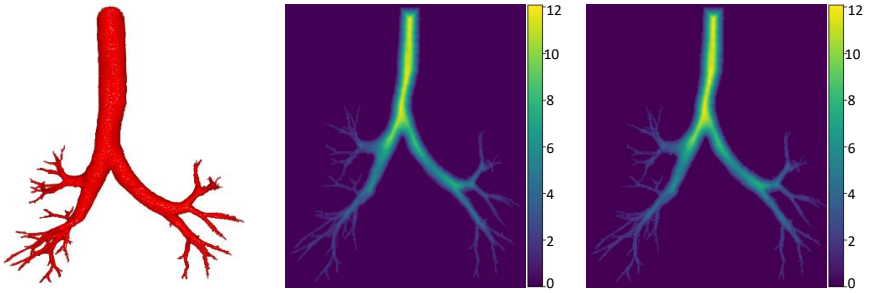

(a) (b) (c)

Fig. 2: Examples of different DT maps.: **(a)** Ground truth, **(b)** Traditional distance transform, **(c)** Soft distance transform. For clear visualization, we provide the max intensity projection maps of DT maps.

and L_{reg} for DT regression. For these, we apply Dice loss and mean square error loss, respectively. As for the unsupervised training, according to Eqs. 6 and 8, the proposed cross geometry consistency and Hausdorff distance consistency could be applied to both the labeled data and unlabeled data in an unsupervised way. The overall loss function of the semi-supervised training method is written as

$$\begin{aligned} L_{overall}(\mathbf{P}_l, \mathbf{P}, \mathbf{G}, \mathbf{Q}, \mathbf{R}) = L_{seg}(\mathbf{P}_l, \mathbf{G}) + L_{reg}(\mathbf{R}, \mathbf{Q}) \\ +\alpha(L_{cross}(\mathbf{P}, \mathbf{R}) + L_{hdc}(\mathbf{P}, \mathbf{R})), \end{aligned} \tag{9}$$

where $\mathbf{P}_l$, $\mathbf{G}$, and $\mathbf{Q}$ denote the predicted segmentation maps of labeled data, the corresponding ground truth, and the corresponding DT ground truth, respectively. α is a trade-off parameter.

3 Experiments and results

To assess our proposed method, we conducted experiments on bronchus segmentation using a binary airway segmentation data set [18]. The data set consists of 90 CT scans (70 from LIDC [1] and 20 from the training set of the EXACT'09 [11]) with slice thicknesses between 0.5 and 1.0 mm. The CT slices had dimensions of 512×512 pixels, with a resolution ranging from 0.50 to 0.82 mm. We randomly selected 80 cases for training and the remaining 10 cases for testing. We employed the tree length detected rate (TD) [4], branch detected rate (BD) [4], and precision rate as our evaluation metrics. Specifically, TD is the percentage of ground truth skeleton voxels within predictions relative to the total ground truth skeleton length. BD measures the predicted branches as a percentage of the actual ground truth branches.

In the training phase, we used an SGD optimizer with an initial learning rate of 0.01, decaying by 0.1 every 2000 iterations. The batches comprised 4 images, split equally between labeled and unlabeled, and we employed random cropping

Table 1: Evaluation results over the 8-protocol and 16-protocol settings.

Approach	Labeled/Unlabeled	TD (%)	BD (%)	Precision (%)
3D U-Net [3]	8/0	80.7±2.5	72.1±2.9	93.6±0.5
3D U-Net [3]	16/0	82.8±0.2	73.5±0.5	92.2±0.4
3D U-Net [3]	80/0	88.8±3.9	89.6±2.6	91.3±0.2
MT [20]	8/72	80.4±0.1	72.5±0.1	93.4±0.2
SASSNet [8]	8/72	79.3±0.9	71.2±1.4	93.5±0.2
DTC [12]	8/72	80.2±1.1	71.8±0.8	93.0±0.9
ICT [21]	8/72	80.9±1.4	71.4±2.2	93.7±1.0
Co-BioNet [16]	8/72	80.1±0.4	70.6±0.4	**94.9±0.1**
Our method	8/72	**87.6±0.1**	**83.0±0.2**	88.8±0.1
MT [20]	16/64	83.1±0.1	75.3±0.4	92.2±0.1
SASSNet [8]	16/64	81.5±1.1	72.5±1.7	**93.2±0.2**
DTC [12]	16/64	81.7±1.0	73.3±1.2	92.8±0.6
ICT [21]	16/64	80.7±3.6	72.4±4.9	92.9±1.0
Co-BioNet [16]	16/64	82.9±3.1	73.5±4.3	91.9±0.9
Our method	16/64	**88.4±0.5**	**82.8±0.5**	89.7±0.8

(size $112 \times 112 \times 80$ voxels), rotation, and flipping for data augmentation. The trade-off parameter λ was adjusted via a time-dependent Gaussian warming-up approach [20, 24]. We conducted two experiments with 8-protocol and 16-protocol settings (i.e., the numbers of labeled training samples were 8 and 16).

We compared our method with five state-of-the-are methods, including MT [20], SASSNet [8], DTC [12], ICT [21], and Co-BioNet [16]. The previous methods were reproduced using the 3D U-Net as the backbone architecture, as same as our method. The experimental results are shown in Table 1. In the 8-protocol setting, our approach surpassed comparison methods with the highest TD and BD rates (TD: 87.6%, BD: 83.0%, precision: 88.8%). With the 16-protocol setting, it again led with the highest TD and BD rates (TD: 88.4%, BD: 82.8%, precision: 89.7%), comparable to a fully supervised 3D U-Net's TD performance. However, the precision rate of our method was lower compared to other methods. Several previous methods underperformed 3D U-Net in terms of the TD or BD, which might be caused by the modifications in the segmentation task and model architecture. The qualitative results are shown in Fig. 3. In the emphasized dotted boxes, we observed that our method segmented more small branches, compared to previous methods. We also observed that our method over-segmented several airway branches, causing a false positive issue and decreasing the precision rate. We conducted an ablation study, with 3D U-Net as the baseline, to assess the impact of cross geometry and Hausdorff consistencies. Results, detailed in Table 2, show that both consistencies improved TD and BD metrics. Integrating these consistencies further enhanced the segmentation performance of our method.

Table 2: Ablation study over the 8-protocol and 16-protocol settings.

Approach	Labeled/Unlabeled	TD (%)	BD (%)	Precision (%)
Baseline	8/0	80.7±2.5	72.1±2.9	93.6±0.5
Baseline	16/0	82.8±0.2	73.5±0.5	92.2±0.4
Baseline	80/0	88.8±3.9	89.6±2.6	91.3±0.2
Baseline + cross	8/72	82.0±0.4	74.2±1.3	**95.2±0.2**
Baseline + Hausdorff	8/72	82.1±1.6	72.4±2.1	93.3±0.1
Our method	8/72	**87.6±0.1**	**83.0±0.2**	88.8±0.1
Baseline + cross	16/64	83.7±2.1	72.5±6.0	84.7±1.1
Baseline + Hausdorff	16/64	84.0±1.3	74.4±4.9	**89.7±0.1**
Our method	16/64	**88.4±0.5**	**82.8±0.5**	89.7±0.8

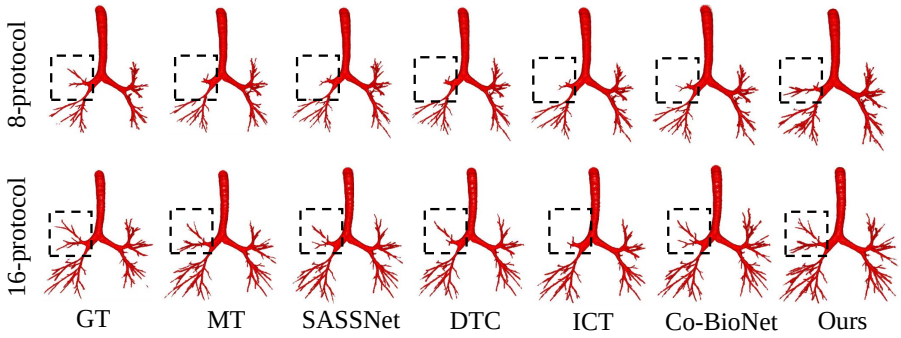

Fig. 3: Segmentation results over 8-protocol and 16-protocol settings.

The precision rate of our method underperformed previous methods and baselines due to increased over-segmentation. This issue arised from two main factors: 1) incomplete annotations: The proposed method detected additional unannotated branches, uncovering incomplete annotations in the ground-truth. This capability led to more comprehensive segmentation of branches not labeled in the ground-truth, owing to the strength of extracting geometric shape features, 2) boundary over-segmentation: The proposed method over-segmented branch boundaries, resulting in a bit thicker branches and false positives, likely due to errors in approximated differentiable distance transform. Combining cross-geometric shape consistency and Hausdorff consistency further extracted geometric shape information, compared to only using one of them. Owing to the combination, the proposed method segmented unannotated airways more precisely, but also introduced some false positives, deteriorating evaluation metrics.

4 Conclusions

In this paper, we present a consistency learning-based semi-supervised method for 3D tubular structure segmentation. We proposed a cross geometry consistency and a Hausdorff distance consistency which could effectively leverage geometry shape and boundary information from unlabeled data. Our method improved the segmentation performance in terms of the tree length rate and branch detected rate. However, the false positive issue exists in our method, decreasing the precision rate of our method. On the other hand, we evaluated our method only on the airway segmentation task in this paper. Our method theoretically works on other tubular structure tasks such as the blood vessel segmentation task. The improvements in precision rate and evaluations on other tubular structure segmentation tasks will be our future work.

Acknowledgements

This work was funded by grants from the JSPS KAKENHI (24H00720, 24K03262), the JST CREST (JPMJCR20D5), the JST [Moonshot R&D] (JPMJMS2033, JPMJMS2214), and the JSPS Bilateral Joint Research Project.

Disclosure of Interests

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

References

1. Armato III, S.G., McLennan, G., Bidaut, L., McNitt-Gray, M.F., Meyer, C.R., Reeves, A.P., Zhao, B., Aberle, D.R., Henschke, C.I., Hoffman, E.A., et al.: The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. Medical Physics **38**(2), 915–931 (2011)
2. Audebert, N., Boulch, A., Le Saux, B., Lefèvre, S.: Distance transform regression for spatially-aware deep semantic segmentation. Computer Vision and Image Understanding **189**, 102809 (2019)
3. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. LNCS. vol. 9901, pp. 424–432. Springer (2016)
4. Feuerstein, M., Kitasaka, T., Mori, K.: Adaptive branch tracing and image sharpening for airway tree extraction in 3-D chest CT. In: Proceedings of Second International Workshop on Pulmonary Image Analysis. vol. 1, pp. 1–8 (2009)
5. French, G., Laine, S., Aila, T., Mackiewicz, M., Finlayson, G.: Semi-supervised semantic segmentation needs strong, varied perturbations. arXiv preprint arXiv:1906.01916 (2019)

6. Guan, D., Huang, J., Xiao, A., Lu, S.: Unbiased subclass regularization for semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9968–9978 (2022)
7. Karimi, D., Salcudean, S.E.: Reducing the hausdorff distance in medical image segmentation with convolutional neural networks. IEEE Transactions on Medical Imaging **39**(2), 499–513 (2019)
8. Li, S., Zhang, C., He, X.: Shape-aware semi-supervised 3D semantic segmentation for medical images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. LNCS. pp. 552–561. Springer (2020)
9. Liu, S., Li, Y., Li, X., Cao, G.: Shape-aware multi-task learning for semi-supervised 3D medical image segmentation. In: 2021 IEEE International Conference on Bioinformatics and Biomedicine. pp. 1418–1423. IEEE (2021)
10. Liu, Z., Zhao, C.: Semi-supervised medical image segmentation via geometry-aware consistency training. arXiv preprint arXiv:2202.06104 (2022)
11. Lo, P., Van Ginneken, B., Reinhardt, J.M., Yavarna, T., De Jong, P.A., Irving, B., Fetita, C., Ortner, M., Pinho, R., Sijbers, J., et al.: Extraction of airways from CT (EXACT'09). IEEE Transactions on Medical Imaging **31**(11), 2093–2107 (2012)
12. Luo, X., Chen, J., Song, T., Wang, G.: Semi-supervised medical image segmentation through dual-task consistency. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 8801–8809 (2021)
13. Ma, J., Wei, Z., Zhang, Y., Wang, Y., Lv, R., Zhu, C., Gaoxiang, C., Liu, J., Peng, C., Wang, L., et al.: How distance transform maps boost segmentation CNNs: an empirical study. In: Medical Imaging with Deep Learning. pp. 479–492. PMLR (2020)
14. Navarro, F., Shit, S., Ezhov, I., Paetzold, J., Gafita, A., Peeken, J.C., Combs, S.E., Menze, B.H.: Shape-aware complementary-task learning for multi-organ segmentation. In: International Workshop on Machine Learning in Medical Imaging. pp. 620–627. Springer (2019)
15. Ouali, Y., Hudelot, C., Tami, M.: Semi-supervised semantic segmentation with cross-consistency training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12674–12684 (2020)
16. Peiris, H., Hayat, M., Chen, Z., Egan, G., Harandi, M.: Uncertainty-guided dual-views for semi-supervised volumetric medical image segmentation. Nature Machine Intelligence **5**(7), 724–738 (2023)
17. Pham, D.D., Dovletov, G., Pauli, J.: A differentiable convolutional distance transform layer for improved image segmentation. In: Proceedings of the DAGM German Conference on Pattern Recognition. LNCS. pp. 432–444. Springer (2021)
18. Qin, Y., Chen, M., Zheng, H., Gu, Y., Shen, M., Yang, J., Huang, X., Zhu, Y.M., Yang, G.Z.: AirwayNet: a voxel-connectivity aware approach for accurate airway segmentation using convolutional neural networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. LNCS. vol. 11769, pp. 212–220. Springer (2019)
19. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, J.P., Bauer, U., Menze, B.H.: clDice-a novel topology-preserving loss function for tubular structure segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16560–16569 (2021)
20. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. Advances in neural information processing systems **30** (2017)

21. Verma, V., Kawaguchi, K., Lamb, A., Kannala, J., Solin, A., Bengio, Y., Lopez-Paz, D.: Interpolation consistency training for semi-supervised learning. Neural Networks **145**, 90–106 (2022)
22. Wang, C., Hayashi, Y., Oda, M., Itoh, H., Kitasaka, T., Frangi, A.F., Mori, K.: Tubular structure segmentation using spatial fully connected network with radial distance loss for 3D medical images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. LNCS. pp. 348–356. Springer (2019)
23. Wang, Y., Wei, X., Liu, F., Chen, J., Zhou, Y., Shen, W., Fishman, E.K., Yuille, A.L.: Deep distance transform for tubular structure segmentation in CT scans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3833–3842 (2020)
24. Yu, L., Wang, S., Li, X., Fu, C.W., Heng, P.A.: Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. LNCS. vol. 11765, pp. 605–613. Springer (2019)
25. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: CutMix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6023–6032 (2019)
26. Zeng, L.L., Gao, K., Hu, D., Feng, Z., Hou, C., Rong, P., Wang, W.: Ss-tbn: A semi-supervised tri-branch network for covid-19 screening and lesion segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023)
27. Zhang, M., Yang, G.Z., Gu, Y.: Differentiable topology-preserved distance transform for pulmonary airway segmentation. arXiv preprint arXiv:2209.08355 (2022)