
clDice - a Novel Connectivity-Preserving Loss Function for Vessel Segmentation

Johannes C. Paetzold*
Technical University Munich
johannes.paetzold@tum.de

Suprosanna Shit*†
Technical University Munich
suprosanna.shit@tum.de

Ivan Ezhov
Technical University Munich
ivan.ezhov@tum.de

Giles Tetteh
Technical University Munich
giles.tetteh@tum.de

Ali Ertürk
Helmholtz Zentrum Munich
erturk@helmholtz-muenchen.de

Bjoern Menze
Technical University Munich
bjoern.menze@tum.de

Abstract

Accurate segmentation of vascular structures is an emerging research topic with relevance to clinical and biological research. The connectedness of the segmented vessels is often the most significant property for many applications such as disease modeling for neurodegeneration and stroke. We introduce a novel metric namely *clDice*, which is calculated on the intersection of centerlines and volumes as opposed to the traditional dice, which is calculated on volumes only. Firstly, we tested state-of-the-art vessel segmentation networks using the proposed metric as evaluation criteria and show that it captures vascular network properties superior to traditional metrics, such as the dice-coefficient. Secondly, we propose a differentiable form of *clDice* as a loss function for vessel segmentation. We find that training on *clDice* leads to segmentation with more accurate connectivity information, higher graph similarity and often superior volumetric scores.

1 Introduction

Segmentation of blood vessels is a key step in many clinical and biological applications such as analyzing neurodegenerative diseases, e.g. Alzheimer’s disease [1], brain-vessel, and stroke modeling [2]. The two most commonly used categories of quantitative performance measures are a) overlap based distance measures such as dice-score, precision, recall, and Jaccard index; and b) volumetric distance measures such as the Hausdorff distance and the Mahalanobis distance [3, 4, 5, 6, 7].

However, in many vessel segmentation applications, the most important properties are the connectivity of the vascular network segments. Traditional scores, e.g. dice and Jaccard rely on the average voxel-wise hit or miss prediction [8]. On the other hand, in a task like vascular network extraction, a proper sequence of hits in the voxel domain is preferred over spurious hits. Further, a globally averaged metric does not equally weight vessels with large, medium and small radii. In real datasets, where vessels of wide radius ranges exist, e.g. $30\ \mu\text{m}$ for arterioles [6, 9] and $5\ \mu\text{m}$ for capillaries, training on a globally averaged loss induces a strong bias towards the volumetric segmentation of

* Authors contributed equally and should be considered as joint first authors.

† Corresponding Author.




Figure 1: **Motivation**, (a) Shows an exemplary 2D slice of real microscopic data, (b) and (c) are two random segmentation results which achieve similar dice scores (not from our presented model). Note that (b) does not capture any of the small vessels while segmenting the large vessel very accurately, on the other side segmentation (c) captures all vessels in the image while being a less accurate on the diameter of the large vessel.

large vessels. This is pronounced in imaging modalities like fluorescence microscopy, where the image intensities of arteriole surrounding tissue are higher than the intensity within capillaries [6]. In Figure 1, an example illustrates the suboptimality of traditional scores in some scenarios.

Furthermore, the most traditional metrics are ambiguous when some of the objects of interest are of the same order as the resolution of the signal. Single-voxel shifts in prediction change the local metric score significantly, thus making the metric difficult to interpret [8]. In this context of a vascular network extraction task, we ask the following research questions:

- Q1. What is a good connectivity-aware metric to benchmark vessel-segmentation algorithms?
- Q2. How can we leverage this metric in a loss function to improve connectivity in vessel-segmentation?

2 Methods

In this section we first introduce the *clDice* as a metric and subsequently introduce a differentiable loss function namely *soft-clDice*.

***clDice* Metric** : We propose a novel connectivity-preserving metric to evaluate vessel segmentation, based on intersecting centerlines of vessels with vessel volumes. We call this metric a centerline-in-volume-dice-coefficient or *clDice* in short. We consider two binary volumes: first, the ground truth label map (L), and second, the predicted segmentation volume (P). The centerlines cl_P and cl_L are extracted from P and L respectively. Subsequently, we compute the fraction of cl_L that lies within P , which we call cl_L2vol_P and vice-a-versa to get cl_P2vol_L [c.f. Algorithm 1]. We observe that cl_P2vol_L is very susceptible to false positives in the prediction while cl_L2vol_P is susceptible to false negatives. Therefore, we proceed to interpret cl_L2vol_P as precision and cl_P2vol_L as recall. Since we want to maximize both precision and recall, we formulate it symmetrically similar to the dice coefficient in Equation 1. This leads us to the final expression of *clDice* in Equation 2.

$$Dice = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}} \quad (1) \quad clDice = 2 \times \frac{cl_P2vol_L \times cl_L2vol_P}{cl_P2vol_L + cl_L2vol_P} \quad (2)$$

***soft-clDice* Loss** : The centerline can be extracted through Euclidean distance transform or via repeated morphological thinning. Although Euclidean distance transform has been used in multiple occasions [10] to induce skeletons, it is a discrete operation and an end-to-end differentiable approximation remains unsolved, which prevents us from using it in our loss function. On the contrary, morphological thinning consists of dilation and erosion operations. Further, min- and max filters are commonly used as the greyscale alternative of morphological dilation and erosion. Motivated by this fact we replace dilation and erosion operations with iterative min- and max-pooling. This allows us to leverage *clDice* to extract a parameter-free, morphologically motivated soft-centerline on greyscale valued data. We call this loss *soft-clDice* and describe it in Algorithm 1 and 2. We determine the hyper-parameter k to be in the range of the maximum radius for




Figure 2: 2D slice of the soft-centerline(right) of a real valued class probability map (left).

the vessel like structure. In our experiment it is 5 for the synthetic and real data. Choosing a larger k does not reduce performance but increases computation time, on the other hand a too low k leads to incomplete skeletonization.

Algorithm 1: *soft-clDice*

Input : P, L
 $cl_P \leftarrow \text{soft-centerline}(P)$
 $cl_L \leftarrow \text{soft-centerline}(L)$
 $cl_P 2vol_L \leftarrow \frac{|cl_P \circ L| + \epsilon}{|cl_P| + \epsilon}$
 $cl_L 2vol_P \leftarrow \frac{|cl_L \circ P| + \epsilon}{|cl_L| + \epsilon}$
 $clDice \leftarrow 2 \times \frac{cl_P 2vol_L \times cl_L 2vol_P}{cl_P 2vol_L + cl_L 2vol_P}$

Output : $clDice$

Algorithm 2: *soft-centerline*

Input : I, k
 $I' \leftarrow \text{maxpool}(\text{minpool}(I))$
 $cl \leftarrow \text{ReLU}(I - I')$
for $i \leftarrow 0$ **to** k **do**
 $I \leftarrow \text{minpool}(I)$
 $I' \leftarrow \text{maxpool}(\text{minpool}(I))$
 $cl \leftarrow cl + cl \circ \text{ReLU}(I - I')$
end
Output : cl

Vessel Segmentation: We evaluate the proposed *clDice* metric on two state-of-the-art 3D vessel segmentation networks i) a 3D U-net[11], and ii) a 3D fully connected network (FCN)[12]. We used generalized *soft-Dice* [13, 14] to train our baseline model for the vessel segmentation. Since our objective here is to achieve accurate segmentation, while giving vascular connection more importance, we add our proposed *soft-clDice* with *soft-Dice* as following

$$\mathcal{L}_e = 0.5(\text{soft-Dice} + \text{soft-clDice}) \quad (3)$$

In stark contrast to previous works, where vessel segmentation and centerline prediction has been learned jointly as a multi-task learning [15], or which only learned vessel-centerlines (or trees [12]), we are not interested in learning the centerlines. We are interested in learning a whole vessel segmentation, where the connections between individual vessels are robust and complete.

3 Experiments

Dataset: We test our proposed metric and loss function on a synthetic and a real dataset. The generation of the synthetic data is described in [16], additionally, we add a Gaussian noise term to this generated data. The real dataset consists of multi-channel volumetric scans of the brain vasculature (voxel size: $(3\mu m^3)$), which were obtained using light-sheet microscopy of tissue cleared Murine brains, as introduced in [17]. We performed experiments on a synthetic dataset using fifteen single-channel volumes for training, two for validation and five for testing, each of the size $325 \times 304 \times 600$ pixels. On the real data we used both single and two-channel inputs, the inputs correspond to different fluorescent stains, which have been shown to contain complimentary information [6]. Eleven volumes were used for training, two for validation and four for testing, each of the size $500 \times 500 \times 50$ pixels.

Evaluation Metric: We report overlap based metrics such as the Dice coefficient, Jaccard index (IOU) and Accuracy along with our proposed *clDice* for all the experiment settings. Additionally, we extract a vascular graph from the centerline of the predicted segmentation and compute relative accuracy of total vascular network length (Dist.), the number of detected bifurcation points (Bifurc.) and endpoints (End Pt.) compared to the ground truth.

Results & Discussion: We trained a Unet and a FCN in different scenarios of identical settings and datasets. From Table 1 we observe that the inclusion of *soft-clDice* loss not only leads to a higher *clDice* in all cases, but also performs better than the standalone *soft-dice* in terms of dice coefficient and IOU. We also observe that *soft-clDice* improves the extracted network properties significantly for real data. We do not see any systematic change in synthetic data after adding *soft-clDice*. We attribute this to the fact that the synthetic data has higher signal-to-noise ratio and lacks significant illumination variation.

4 Conclusion

This abstract introduces a novel connectivity-preserving metric *clDice* for vessel segmentation. We use the new metric to evaluate segmentation quality, and in a loss function, to train state-of-the-art networks on real and synthetic data. We find that training on *soft-clDice* leads to vessel segmentation

Table 1: Experimental results for 3D U-nets and 3D FCNs on synthetic and real data. We observe a consistent performance improvement for real data with the combination of *soft-clDice* and *soft-dice*.

Data	Loss	Network	Dice	clDice	IOU	Acc.	Dist.	Bifurc.	End Pt.
Synthetic	<i>soft-dice</i>	FCN, 1 ch	99.41	99.45	98.83	99.97	0.92	0.91	0.91
		Unet, 1 ch	99.61	99.90	97.23	99.98	0.88	0.86	0.89
	\mathcal{L}_c	FCN, 1 ch	99.16	99.77	98.34	99.96	0.92	0.91	0.92
		Unet, 1 ch	98.73	99.90	97.49	99.94	0.88	0.86	0.88
Real data	<i>soft-dice</i>	FCN, 1 ch	75.28	90.98	60.35	89.88	0.87	0.72	0.81
		FCN, 2 ch	78.54	92.03	64.67	91.66	0.90	0.82	0.84
		Unet, 1 ch	87.11	95.03	77.17	95.78	0.92	0.82	0.97
		Unet, 2 ch	80.20	93.05	66.94	92.33	0.95	0.93	0.70
	\mathcal{L}_c	FCN, 1 ch	85.57	96.16	74.78	95.09	0.97	0.88	0.97
		FCN, 2 ch	85.28	95.75	74.34	94.91	0.91	0.91	0.97
		Unet, 1 ch	86.94	95.28	76.89	95.86	0.94	0.83	0.97
		Unet, 2 ch	83.96	96.10	72.36	94.18	0.96	0.89	0.85

with more accurate connectivity information, higher graph similarity and similar to better volumetric scores. More importantly *clDice* and *soft-clDice* can be readily deployed in other tree-structured object segmentation tasks such as neuron segmentation and bronchial tract segmentation.

Acknowledgments

Suprosanna Shit and Ivan Ezhov are supported by the Translational Brain Imaging Training Network (TRABIT) under the European Union’s ‘Horizon 2020’ research & innovation program (Grant agreement ID: 765148). The authors would like to thank the Deutsche Forschungsgemeinschaft (DFG; ME3511/3-1). We thank Mihail I. Todorov, Amirhossein Bayat and Oliver Schoppe.

References

- [1] Jesse M Hunter et al. Morphological and pathological evolution of the brain microcirculation in aging and Alzheimer’s disease. *PLoS one*, 7(5):e36893, 2012.
- [2] Anne Joutel et al. Cerebrovascular dysfunction and microcirculation rarefaction precede white matter lesions in a mouse genetic model of cerebral ischemic small vessel disease. *The Journal of Clinical Investigation*, 120(2):433–445, 2010.
- [3] Cemil Kirbas and Francis Quek. A review of vessel extraction techniques and algorithms. *ACM Computing Surveys (CSUR)*, 36(2):81–121, 2004.
- [4] Matthias Schneider et al. Joint 3-D vessel segmentation and centerline extraction using oblique Hough forests with steerable filters. *Medical Image Analysis*, 19(1):220–249, 2015.
- [5] Renzo Phellan et al. Vascular segmentation in tof mra images of the brain using a deep convolutional neural network. In *MICCAI Workshop*, pages 39–46. Springer, 2017.
- [6] Mihail Ivilinov Todorov et al. Automated analysis of whole brain vasculature using machine learning. *bioRxiv*, page 613257, 2019.
- [7] Kai Hu et al. Retinal vessel segmentation of color fundus images using multiscale convolutional neural network with an improved cross-entropy loss function. *Neurocomputing*, 309:179–191, 2018.
- [8] Abdel Aziz Taha and Allan Hanbury. Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool. *BMC Medical Imaging*, 15(1):29, 2015.
- [9] Antonino Paolo Di Giovanna et al. Whole-brain vasculature reconstruction at the single capillary level. *Scientific reports*, 8(1):12573, 2018.
- [10] Fernando Navarro et al. Shape-aware complementary-task learning for multi-organ segmentation. In *International Workshop on MLMI*, pages 620–627. Springer, 2019.

- [11] Özgün Çiçek and Aothers. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, pages 424–432. Springer, 2016.
- [12] Giles Tetteh et al. Deepvesselnet: Vessel segmentation, centerline prediction, and bifurcation detection in 3-d angiographic volumes. *arXiv preprint arXiv:1803.09340*, 2018.
- [13] Fausto Milletari et al. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, pages 565–571. IEEE, 2016.
- [14] Carole H Sudre et al. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *MICCAI Workshop*, pages 240–248. Springer, 2017.
- [15] Fatmatülzehra Uslu and Anil Anthony Bharath. A multi-task network to detect junctions in retinal vasculature. In *MICCAI*, pages 92–100. Springer, 2018.
- [16] Matthias Schneider et al. Tissue metabolism driven arterial tree generation. *Medical Image Analysis*, 16(7):1397–1414, 2012.
- [17] Ali Ertürk et al. Three-dimensional imaging of solvent-cleared organs using 3DISCO. *Nature protocols*, 7(11):1983, 2012.