Regional Hausdorff Distance Losses for Medical Image Segmentation

Lisa Guzzi[1,2,3], Maria A. Zuluaga[2], Riccardo Taiello[4], Fabien Lareyre[5,6], Gilles Di Lorenzo[5], Sébastien Goffart[1,3], Andrea Chierici[3], Juliette Raffort[3,6], and Hervé Delingette[1]

[1] Université Côte d'Azur, Inria, Epione Team, Sophia Antipolis, France
[2] Data Science Department, EURECOM, Sophia Antipolis, France
[3] University Hospital of Nice, Nice, France
[4] Independent researcher, Geneva, Switzerland
[5] Department of Vascular Surgery, Hospital of Antibes Juan-les-Pins, Antibes, France
[6] Université Côte d'Azur, CNRS, UMR7370, LP2M, Nice, France
lisa.guzzi@inria.fr

Abstract. Most medical image segmentation techniques rely on overlap-based metrics such as the Dice coefficient. In contrast, the Hausdorff Distance (HD) offers a more sensitive assessment of boundary discrepancies by explicitly capturing spatial misalignments. Despite its relevance, directly minimizing the HD during the training of convolutional neural networks for medical image segmentation remains challenging due to the non-differentiability of the conventional distance transform algorithms. Previous attempts of soft distance transforms are limited by numerical instability or require binary inputs, limiting their applicability. In this paper, we introduce novel regional Hausdorff Distance loss functions to optimize the HD without relying on any auxiliary losses. Specifically, we propose the maximum, modified, and average regional Hausdorff Distance losses. Central to our approach is a new method to compute a fully differentiable erosion-based distance function, which can be applied directly to probability maps. These functions accurately approximate the signed, unsigned, or positive distance maps while maintaining full differentiability. We validate our approach on multiple public medical image segmentation datasets, demonstrating that our HD losses achieve competitive performance, outperforming state-of-the-art methods.

Keywords: Hausdorff Distance · Image Segmentation · Distance transforms.

1 Introduction

The Hausdorff Distance (HD) metric quantifies the distance between two point sets [11]. In medical image analysis, it is widely used to assess the segmentation performance of deep learning models [6, 24]. Unlike the Dice Similarity Coefficient (DSC), which only measures the volumetric overlap between a predicted and the ground truth segmentation, the HD metric offers a more sensitive

assessment of boundary discrepancies by explicitly capturing the maximal distance between misaligned parts. This makes it particularly relevant for assessing the quality of the delineation of small structures such as tumors or lesions by highlighting the worst-case mismatch between two objects [19], or the boundary alignment of anatomical structures like the heart or the liver [15]. Despite its importance, few attempts have been reported to directly optimize the HD metric during the training of deep learning-based image segmentation. Some studies have designed losses to optimize the boundary accuracy without explicitly reproducing the HD [14, 5, 27, 9]. The weighted Hausdorff Distance loss from [18] approximates the average HD for isolated points but is designed only for points, not image masks. Only the work from [13] aims to reduce the HD by approximating an HD loss, relying on the Euclidean Distance Transform (EDT) of the segmentation boundaries. They also propose an alternative loss through morphological operations, which underperformed compared to the EDT-based formulation. While this loss improved performance regarding HD, it must be combined with Dice loss to prevent any instabilities, and its formulation does not exactly replicate the HD formula. Additionally, the EDT-based loss lacks differentiability properties. Indeed, Distance Transform Maps (DTM) provide an alternative representation of a binary shape where each voxel's intensity corresponds to its distance with the nearest foreground boundary voxel [21]. This representation enables explicit distance computations between point sets. However, conventional distance transform algorithms are inherently non-differentiable [28]. The differentiability of the approach in [13] is not explicitly addressed, raising questions about its suitability for gradient-based optimization. Several studies have tried approximating soft distance transforms (DT) by emulating the operation through CNN-based learning [2, 4, 16]. For instance, in [2], a CNN-based approach was introduced to learn a watershed transform for instance-based segmentation. Other methods have directly regressed the DTM as a CNN output to be used as a loss function or as a regularization term [25, 7, 26]. However, these approaches are prone to domain shift since the learned DT is dataset-dependent, and some require additional geometry-aware refinements to improve generalization. Recent works have proposed Convolutional-based differentiable Distance Transform (CDT) operations [29, 28, 17]. One major limitation of CDT is that the kernel size must be as large as the image diagonal to process sparse binary images correctly. If the kernel size is too small, background pixels far from the foreground will mistakenly receive a zero-distance value. This leads to increased computational complexity due to the large kernel size and potential numerical instability when the kernel's exponential term approaches zero, particularly in large images with only a few foreground pixels. A cascaded CDT approach was proposed by [17] to mitigate these issues, but it requires soft binarization of segmentation outputs, hence not directly applicable to raw probability maps.

In this paper, we propose a novel family of regional HD-based loss functions that rely on a soft and differentiable DT. As a first contribution we introduce a novel morphological erosion-based differentiable distance transform that can be applied directly on probability maps. This method allows to compute the ac-

curate signed, unsigned, and positive distance maps in a differentiable manner, ensuring compatibility with deep learning models. In a second contribution, we derive three well-established variations of the HD loss leveraging the differentiable distance transform. Specifically, we provide a smooth formulation of the Hausdorff, Modified Hausdorff, and Symmetric Averaged Hausdorff Distances. These losses achieve state-of-the-art performance without auxiliary losses.

We analyze the effect of key hyperparameters and validate our method through controlled validation experiments. Then, we evaluate the applicability of our losses across various public medical segmentation datasets, proving the applicability of the soft distance transforms and the effectiveness of the proposed loss functions to minimize the distance of the predicted masks to the ground truth while maintaining an equivalent Dice score. To foster reproducibility, our code is available at https://github.com/lisaGUZZI/HD-Loss

2 Method

2.1 Definition of distance transforms

We define a ground truth binary segmentation G consisting of N voxels such that $G = \{G_n\}_{n=1}^{N}, \quad G_n \in \{0, 1\}$. The boundary ∂G of the foreground object in the image is defined as $\partial G = \{q \mid G_q = 1,\ \exists r \text{ such that } G_r = 0,\ dt(q, r) \leq 1\}$ where $dt(q, r)$ is a chosen distance metric between two voxels, typically the Euclidean or Manhattan distance.

The distance transform map $\mathcal{D}(G)$ assigns to each voxel G_n its minimum distance to the closest foreground boundary voxel q such that:

$$\mathcal{D}(G_n) = \min_{q \in \partial G} dt(G_n, q) \tag{1}$$

Eq. 1 actually corresponds to the unsigned distance transform $\mathcal{D}^u$ where the distance is computed for each voxel, both inside and outside the boundary ∂G. Alternatively, one can also compute the signed distance transform $\mathcal{D}^s$ such that the distance is negative inside the object and positive outside:

$$\mathcal{D}^s(G_n) = \begin{cases} \min_{q \in \partial G} dt(G_n, q), & \text{if } G_n = 0 \quad \text{(outside the object)} \\ -\min_{q \in \partial G} dt(G_n, q), & \text{if } G_n = 1 \quad \text{(inside the object)} \end{cases} \tag{2}$$

The positive distance transform $\mathcal{D}^+$ assigns a distance of 0 to all foreground voxels and computes the distance only for background voxels:

$$\mathcal{D}^+(G_n) = \begin{cases} \min_{q \in \partial G} dt(G_n, q), & \text{if } G_n = 0 \\ 0, & \text{if } G_n = 1 \end{cases} \tag{3}$$

The relationships between these distance transforms are therefore given by: $\mathcal{D}^u(G) = \mathcal{D}^+(G) + \mathcal{D}^+(1-G)$ and $\mathcal{D}^s(G) = \mathcal{D}^+(G) - \mathcal{D}^+(1-G)$. Also, $\mathcal{D}^u(G) = |\mathcal{D}^s(G)|$ and $\mathcal{D}^+(G) = \max(\mathcal{D}^s(G), 0)$.

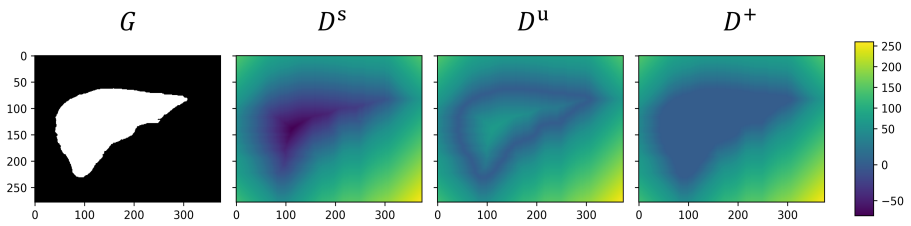

Fig. 1: Illustration of distance transforms $\mathcal{D}^s$, $\mathcal{D}^u$, $\mathcal{D}^+$ on a binary image G.

2.2 Distance transforms with morphological operations

The positive distance transform $\mathcal{D}^+$ can be computed with morphological operations, specifically using successive erosions, when morphological distances are considered. The maximum distance corresponds to the number of erosions required to completely remove the background. The erosion filter $\ominus$ is defined by an isotropic structuring element k (typically $k = 4 + 1$ or $k = 8 + 1$ in 2D, and $k = 6 + 1$, $k = 18 + 1$ or $k = 26 + 1$ in 3D) which defines the unit ball of the considered metric. The erosion of shape G with the k structuring element is written as $G \ominus k$, and we note $G \ominus^i k$ its ith eroded shape.

We propose to compute the positive distance transform $\mathcal{D}^+$ as:

$$\mathcal{D}^+(G) = \sum_{i=1}^{I} [(1 - G) \ominus^i k] \tag{4}$$

where I represents the number of iterations required to compute the complete distance map, or in other words, the number of successive erosions to perform to remove the background. Therefore, each background voxel receives a value equal to the number of erosions required to remove the shape complement $1 - G$. The other distance transforms $\mathcal{D}^u$ and $\mathcal{D}^s$ can then be derived from their respective relationships to $\mathcal{D}^+$.

2.3 Soft distance transforms

Neural networks typically output a segmentation $\mathcal{Y}$ of an image X corresponding to the probability map $\mathcal{Y} = \{y_n\} \in [0, 1]^N$ with $y_n = p(G_n = 1|X) \in [0, 1]$. We want to formalize the $\mathcal{D}^u$, $\mathcal{D}^s$ and $\mathcal{D}^+$ distance transforms to their probabilistic equivalents d^u, d^s and d^+ which can be applied on the probabilistic segmentation $\mathcal{Y}$. The probabilistic distance transform d should generalize to the discrete case $\mathcal{D}(Y) = d(Y)$ such that both give the same output when Y is binary.

Given the soft and differentiable morphological filters introduced in [10], we propose to compute the soft distance transform by replacing the binary erosion filter $\ominus^i k$ in Eq. 4 with a soft erosion filter $\ominus_s^i k$ applied on the probability segmentation $\mathcal{Y}$. The soft erosion $\ominus_s$ is expressed as the multi-linear polynomial

$(\mathcal{Y} \ominus_s k)_n = \prod_{j \in \mathcal{N}^k(n)} \mathcal{Y}_j$, where $\mathcal{N}^k(n)$ are the neighboring pixels of voxel n according to the selected structuring element. This soft erosion operator makes the soft distance transforms applicable on probability maps and fully differentiable.

2.4 Regional Hausdorff Distance Losses

The Hausdorff Distance (HD) quantifies the maximum discrepancy between two point sets. Given the point sets A and B, the HD is defined as: $HD = \max(\max_{a \in A} dt(a, B), \max_{b \in B} dt(b, A))$ where the directed distance of a point a to the set B is given by $dt(a, B) = \min_{b \in B} dt(a, b)$. This formulation of the HD satisfies the triangle inequality property. However, its sensitivity to outliers has led to the introduction of alternative HD formulations. The Modified Hausdorff Distance (MHD) [8] averages the minimal distance instead of taking the maximum. The Symmetric Averaged Hausdorff Distance (MHD_{sym}) [8] takes the mean of the two directed distances. The MHD is more robust in matching objects based on edge points than the HD [8], while MHD_{sym} provides a smooth alternative with a slightly reduced discriminatory power.

To optimize the HD into neural network-based segmentation, we propose differentiable approximations of the HD formulas based on the positive distance transform $\mathcal{D}^+(G)$ and its probabilistic counterpart $d^+(\mathcal{Y})$. Given a binary ground truth segmentation G and a probabilistic prediction segmentation $\mathcal{Y}$, we define: **The Hausdorff Loss (L_H)**, approximating the classical HD:

$$L_H = \text{smax}(\text{smax}(\mathcal{D}^+(G)_n \circ y_n), \text{smax}(d^+(y)_n \circ G_n)) \tag{5}$$

The Averaged Hausdorff Loss (L_{AH}):

$$L_{AH} = \text{smax}\left(\frac{1}{|Y| + \epsilon} \sum_n \mathcal{D}^+(G)_n \circ y_n, \quad \frac{1}{|G| + \epsilon} \sum_n d^+(\mathcal{Y})_n \circ G_n\right) \tag{6}$$

The symmetric Averaged Hausdorff Loss ($L_{AH_{sym}}$):

$$L_{AH_{sym}} = \frac{1}{2}\left(\frac{1}{|Y| + \epsilon} \sum_n \mathcal{D}^+(G)_n \circ y_n + \frac{1}{|G| + \epsilon} \sum_n d^+(\mathcal{Y})_n \circ G_n\right) \tag{7}$$

where $\text{smax}(x)$ is a smooth approximation of the maximum operator implemented using the LogSumExp function (with a α scaling parameter), $\circ$ represents the Hadamard product and ϵ a small positive constant. The directed distance $\mathcal{D}^+(G)_n \circ y_n$ represents the false positive points of $\mathcal{Y}$ and their distance to the ground truth G whereas $d^+(y)_n \circ G_n$ is the directed distance representing the distance of the false negatives to G.

Computational complexity. Considering that the complexity of the soft dice loss is $\mathcal{O}(b \cdot c \cdot V)$, the proposed losses have a complexity of $\mathcal{O}(I \cdot b \cdot c \cdot V)$ where I is the number of soft erosion iterations, and b,c,V denote batch size, number of

channels, and number of voxels, respectively. To reduce computational overhead, we propose constraining I to a small fixed value (e.g., 2–5), resulting in a saturated positive distance transform. The truncation attenuates the influence of distant errors without removing them, maintaining meaningful gradient signals and additionally improving numerical stability during optimization, ensuring that overly large distances do not dominate the loss.

3 Results

3.1 Experimental Setup

Datasets. We evaluated our method on five public datasets: the DRIVE dataset of 2D retinal blood vessels [23], the 3D CT Liver, Spleen and Pancreas datasets from the Medical Segmentation Decathlon [1], and the ACDC dataset [3] comprising multi-slice 2D cine MRI images of the heart. 3D datasets were considered as a stack of 2D slices in the 2D applications.
Evaluation Metrics. We evaluated segmentation performance using the Dice coefficient, conventional HD, HD95, modified HD (MHD) to verify that Hausdorff distance losses regress distance metrics, and clDice [22] to assess the global shape and object-level alignment.
Implementation Details. Experiments were conducted with Python 3.12.8 and Pytorch 2.5.1 on 3 Nvidia A40 PCIe GPUs.

3.2 Validation of Distance transforms

Our differentiable distance transforms, based on morphological operations, is sensitive to the connectivity choice. To evaluate the validity of the generated distance maps, we analyzed the variance of the gradient norm of the signed distance map computed with $\mathcal{D}^s$ under different connectivity settings. Indeed, the ideal configuration should exhibit minimal variance, as the norm of the gradient should be equal to 1 everywhere: $\|\nabla \mathcal{D}^s(G)\| = 1$. We considered signed distance maps produced with a connectivity of 4 and 8 in 2D and of 6, 18, and 26 in 3D. We additionally compared these results with i) the *distance_transform_edt* function from scipy which provides the exact distance but is not differentiable, and ii) the convolutional distance transforms from [29] and [17]. In Table 1, we reported the variance averaged over a database of 50 images randomly selected from the DRIVE, Liver, Spleen, and ACDC datasets, both in 2D and 3D. In 2D, our method achieved a lower norm of the gradient variance around the value 1 than SciPy with both 8 and 4-connectivity. In 3D, the Scipy package got the smallest variances, and the 26-connectivity led to the smallest variance to 1. In the remainder, a connectivity of 4 is selected in 2D and of 26 in 3D.

3.3 Evaluation on Public Datasets

2D. We used our Hausdorff loss functions to train a 2D U-Net [20] for the segmentation of the DRIVE and Spleen datasets (Table 2). Datasets were split

Table 1: Variance of the norm of the gradient of different SDT

Modality	**DT**		**Variance to 1 ↓**	**Variance to mean ↓**
2D	Scipy		0.96	0.03
	CDT [29]		2.80	2.16
	cascaded CDT [17]		0.24	0.21
	Ours	*4 connec*	0.14	**0.02**
		8 connec	**0.12**	0.12
3D	Scipy		**0.97**	**0.03**
	CDT [29]		4.68	4.18
	Ours	*6 connec*	74.64	72.63
		18 connec	21.87	21.62
		26 connec	19.89	19.69

into 75% training, 15% validation, and 15% testing sets, and the network was trained with a batch size of 16 through 250 epochs and a learning rate of 1e-3 for the DRIVE dataset and of 1e-4 for the Spleen. We set the number of erosion to compute the distance transform to 5 iterations for the L_H loss and to 2 for L_{AH} and $L_{AH_{sym}}$. The α scale factor of the LogSumExp function was empirically set to $\alpha = 150$. We compared our method with models trained with the same UNET and the Dice loss as a baseline and the HD loss proposed in [13].

Our loss functions improved the HD metrics (HD, HD95, MHD) compared to the baseline (Dice loss) while maintaining equivalent Dice and clDice scores on both datasets. The Hausdorff loss from [13] did not show any improvement compared to the baseline. When we compared the number of erosion iterations to compute the distance map (c.f. supplementary material), increasing the number of erosion for the L_H loss decreased the HD metrics in both datasets. However, the averaged losses L_{AH} and $L_{AH_{sym}}$ did not exhibit the same behavior as the performance sometimes decreases as the number of iterations increases. This could be because the LAH criterion is a regularized version of LH, which does not verify the triangular inequality and can entail extra local minima. Indeed, the normalization term by the size of the predicted mask $\mathcal{Y}$ can lead to a local minimum for slightly over-segmented predicted masks. Increasing the number of iterations to compute the distance transform might increase the chances of attaining those local minima. Additionally, while the losses significantly improved the HD metrics on the Spleen dataset, the results were less significant regarding the retinal dataset. We believe this is because the DRIVE dataset contains images of retinal blood vessels that are dense and spread through the entire image while the Spleen contains mask that are smaller and centered in the image, meaning that the distance of false positive points could be higher and the losses correct these segmentation discrepancies more effectively.

3D. We used our Hausdorff loss functions to train a 3D nnU-Net [12] for the segmentation of the Pancreas dataset in 3d full resolution (Table 3). This dataset has two labels: Pancreas and Cancer. Datasets were split into 65% training,

Table 2: 2D U-Net segmentation on the Spleen and DRIVE datasets.

		Dice ↑	**HD ↓**	**HD95 ↓**	**MHD ↓**	**clDice ↑**
Spleen	*Dice (baseline)*	**0.93**	10.95	10.10	5.93	**0.96**
	HD [13]	0.88	15.08	13.97	11.56	0.92
	LH	0.90	6.82	5.62	3.51	0.94
	LAH	**0.93**	**5.51**	**4.50**	**0.87**	**0.96**
	LAH sym	0.92	9.66	8.88	5.32	0.95
Drive	*Dice (baseline)*	**0.82**	16.77	2.20	0.50	**0.84**
	HD [13]	0.78	17.80	3.04	0.67	0.81
	LH	0.80	15.58	2.10	0.49	0.82
	LAH	0.80	15.53	1.96	0.48	0.84
	LAH sym	**0.82**	**14.80**	**1.85**	**0.45**	**0.84**

Table 3: 3D nnU-Net segmentation on the Pancreas dataset.

		Dice ↑	**HD ↓**	**HD95 ↓**	**MHD ↓**	**clDice ↑**
Pancreas	*Dice (baseline)*	0.77	25.14	13.37	524.49	0.72
	LH	0.77	25.22	12.22	519.86	0.73
	LAH	0.80	24.65	10.43	282.03	0.75
	LAH sym	**0.82**	**16.94**	**7.10**	**221.80**	**0.79**
Cancer	*Dice (baseline)*	0.43	42.28	29.95	1018.01	0.48
	LH	0.42	23.75	20.79	948.13	0.52
	LAH	0.48	23.29	20.73	867.23	0.53
	LAH sym	**0.54**	**22.09**	**19.23**	**283.82**	**0.60**

15%validation, and 20% testing sets, and the network was trained with a batch size of 12 through 250 epochs and an initial learning rate of 1e-2. We set the number of erosion to compute the distance transform to 2 iterations in every case. The α scale factor of the LogSumExp function was still set to $\alpha = 150$. We compared our methods with an nnU-Net trained with the Dice loss as baseline.

The three Hausdorff losses improved all distance metrics and clDice for all labels especially for the cancer label, excepted for the MHD with L_H that was only slightly improved. The other two losses actually correspond the MHD formulation, explaining why they were more efficient to improve this metric. The Dice score was also improved using L_{AH} compared to the baseline.

4 Conclusion

We introduced regional Hausdorff Distance (HD) losses to regress the Hausdorff, modified Hausdorff, and average Hausdorff Distances in CNN-based medical image segmentation. The introduced method relies on a smooth formulation of the distance transform that can be applied to probability maps. Finally, we have shown that these losses reduce the HD while preserving the Dice score without requiring any auxiliary loss. Furthermore, we show that only two iterated erosions

are sufficient to compute the regional HD losses and improve the performance in terms of HD. In future work, we want to investigate more accurate distance maps in 3D such as Chamfer distances, and the adoption of other smooth maximum functions than the LogSumExp function. This work holds the potential to improve the segmentation of medical structures where the HD is of higher relevance, offering a new paradigm for optimizing segmentation models beyond Dice-centric metrics.

Acknowledgments. This work has been supported by the French government, through the 3IA Côte d'Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002. Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see https://www.grid5000.fr).

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

References

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), 4128 (2022)
2. Bai, M., Urtasun, R.: Deep watershed transform for instance segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5221–5229 (2017)
3. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Gonzalez Ballester, M.A., Sanroma, G., Napel, S., Petersen, S., Tziritas, G., Grinias, E., Khened, M., Kollerathu, V.A., Krishnamurthi, G., Rohé, M.M., Pennec, X., Sermesant, M., Isensee, F., Jäger, P., Maier-Hein, K.H., Full, P.M., Wolf, I., Engelhardt, S., Baumgartner, C.F., Koch, L.M., Wolterink, J.M., Išgum, I., Jang, Y., Hong, Y., Patravali, J., Jain, S., Humbert, O., Jodoin, P.M.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: Is the problem solved? IEEE Transactions on Medical Imaging **37**(11), 2514–2525 (2018)
4. Bui, T.D., Wang, L., Chen, J., Lin, W., Li, G., Shen, D.: Multi-task learning for neonatal brain segmentation using 3d dense-unet with dense attention guided by geodesic distance. In: Domain Adaptation and Representation Transfer and Medical Image Learning with Less Labels and Imperfect Data. pp. 243–251. Springer International Publishing, Cham (2019)
5. Caliva, F., Iriondo, C., Martinez, A.M., Majumdar, S., Pedoia, V.: Distance map loss penalty term for semantic segmentation (2019)
6. Crum, W., Camara, O., Hill, D.: Generalized overlap measures for evaluation and validation in medical image analysis. IEEE Transactions on Medical Imaging **25**(11), 1451–1461 (2006)
7. Dangi, S., Linte, C.A., Yaniv, Z.: A distance map regularized cnn for cardiac cine mr image segmentation. Medical physics **46**(12), 5637–5651 (2019)

8. Dubuisson, M.P., Jain, A.K.: A modified hausdorff distance for object matching. In: Proceedings of 12th international conference on pattern recognition. vol. 1, pp. 566–568. IEEE (1994)
9. EL Jurdi, R., Petitjean, C., Honeine, P., Cheplygina, V., Abdallah, F.: A surprisingly effective perimeter-based loss for medical image segmentation. In: Proceedings of the Fourth Conference on Medical Imaging with Deep Learning. Proceedings of Machine Learning Research, vol. 143, pp. 158–167 (07–09 Jul 2021)
10. Guzzi, L., Zuluaga, M.A., Lareyre, F., Di Lorenzo, G., Goffart, S., Chierici, A., Raffort, J., Delingette, H.: Differentiable soft morphological filters for medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 177–187. Springer (2024)
11. Huttenlocher, D., Klanderman, G., Rucklidge, W.: Comparing images using the hausdorff distance. IEEE Transactions on Pattern Analysis and Machine Intelligence **15**(9), 850–863 (1993)
12. 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)
13. 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)
14. Kervadec, H., Bouchtiba, J., Desrosiers, C., Granger, E., Dolz, J., Ben Ayed, I.: Boundary loss for highly unbalanced segmentation. Medical Image Analysis **67**, 101851 (2021)
15. Moradi, S., Oghli, M.G., Alizadehasl, A., Shiri, I., Oveisi, N., Oveisi, M., Maleki, M., Dhooge, J.: Mfp-unet: A novel deep learning based approach for left ventricle segmentation in echocardiography. Physica Medica **67**, 58–69 (2019)
16. 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: Machine Learning in Medical Imaging. pp. 620–627 (2019)
17. Pham, D.D., Dovletov, G., Pauli, J.: A differentiable convolutional distance transform layer for improved image segmentation. In: Pattern Recognition: 42nd DAGM German Conference, DAGM GCPR 2020, Tübingen, Germany, September 28–October 1, 2020, Proceedings 42. pp. 432–444. Springer (2021)
18. Ribera, J., Güera, D., Chen, Y., Delp, E.J.: Locating objects without bounding boxes (2019)
19. Rizzetto, F., Calderoni, F., De Mattia, C., Defeudis, A., Giannini, V., Mazzetti, S., Vassallo, L., Ghezzi, S., Sartore-Bianchi, A., Marsoni, S., et al.: Impact of inter-reader contouring variability on textural radiomics of colorectal liver metastases. European radiology experimental **4**, 1–12 (2020)
20. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. pp. 234–241. Lecture Notes in Computer Science, Springer International Publishing, Cham (2015)
21. Rosenfeld, A., Pfaltz, J.L.: Distance functions on digital pictures. Pattern recognition **1**(1), 33–61 (1968)
22. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, J.P.W., Bauer, U., Menze, B.H.: clDice - a Novel Topology-Preserving Loss Function for Tubular Structure Segmentation. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16555–16564. IEEE Computer Society (Jun 2021)

23. Staal, J., Abramoff, M., Niemeijer, M., Viergever, M., Van Ginneken, B.: Ridge-Based Vessel Segmentation in Color Images of the Retina. IEEE Transactions on Medical Imaging **23**(4), 501–509 (Apr 2004)
24. Taha, A.A., Hanbury, A.: Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. BMC medical imaging **15**, 1–28 (2015)
25. 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)
26. Xue, Y., Tang, H., Qiao, Z., Gong, G., Yin, Y., Qian, Z., Huang, C., Fan, W., Huang, X.: Shape-aware organ segmentation by predicting signed distance maps. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 12565–12572 (2020)
27. Yang, S., Kweon, J., Kim, Y.H.: Major vessel segmentation on x-ray coronary angiography using deep networks with a novel penalty loss function (2019)
28. Zhang, M., Yang, G.Z., Gu, Y.: Differentiable topology-preserved distance transform for pulmonary airway segmentation. arXiv preprint (2022)
29. Zhu, R., Oda, M., Hayashi, Y., Kitasaka, T., Mori, K.: Semi-supervised tubular structure segmentation with cross geometry and hausdorff distance consistency. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. pp. 612–622. Springer Nature Switzerland, Cham (2024)