Nested star-shaped objects segmentation using diameter annotations

Robin Camarasa [a,b,*], Hoel Kervadec [a,b], M. Eline Kooi [e,f], Jeroen Hendrikse [g],
Paul J. Nederkoorn [h], Daniel Bos [b,c], Marleen de Bruijne [a,b,d]

[a] Biomedical Imaging Group Rotterdam, Erasmus MC, Rotterdam, The Netherlands
[b] Department of Radiology and Nuclear Medicine, Erasmus MC, Rotterdam, The Netherlands
[c] Department of Epidemiology, Erasmus MC, Rotterdam, The Netherlands
[d] Department of Computer Science, University of Copenhagen, Denmark
[e] Department of Radiology and Nuclear Medicine, Maastricht University Medical Center, Maastricht, The Netherlands
[f] Cardiovascular Research Institute Maastricht (CARIM), Maastricht University, Maastricht, The Netherlands
[g] Department of Radiology, University Medical Center Utrecht, Utrecht, The Netherlands
[h] Department of Neurology, Academic Medical Center University of Amsterdam, Amsterdam, The Netherlands

ARTICLE INFO

ABSTRACT

Most current deep learning based approaches for image segmentation require annotations of large datasets, which limits their application in clinical practice. We observe a mismatch between the voxelwise ground-truth that is required to optimize an objective at a voxel level and the commonly used, less time-consuming clinical annotations seeking to characterize the most important information about the patient (diameters, counts, etc.).

In this study, we propose to bridge this gap for the case of multiple nested star-shaped objects (*e.g.*, a blood vessel lumen and its outer wall) by optimizing a deep learning model based on diameter annotations. This is achieved by extracting in a differentiable manner the boundary points of the objects at training time, and by using this extraction during the backpropagation. We evaluate the proposed approach on segmentation of the carotid artery lumen and wall from multisequence MR images, thus reducing the annotation burden to only four annotated landmarks required to measure the diameters in the direction of the vessel's maximum narrowing. Our experiments show that training based on diameter annotations produces state-of-the-art weakly supervised segmentations and performs reasonably compared to full supervision.

We made our code publicly available at https://gitlab.com/radiology/aim/carotid-artery-image-analysis/nested-star-shaped-objects.

1. Introduction

In recent years, deep learning has become ubiquitous in medical image segmentation, outperforming most traditional segmentation methods (Litjens et al., 2017). However, the training of deep learning algorithms typically requires a large amount of annotated data, and therefore substantial manual labor from clinical experts. While both in clinical practice and in clinical research images are being annotated at a large scale, we observe a mismatch between the simpler and less time-consuming type of annotations that are made for these purposes and the (dense) annotations needed to train segmentation networks. In this paper, we propose a method that aims to reconcile the objectives of those two types of annotations, training a network to segment the carotid artery based on diameter annotations that are commonly made to establish the degree of stenosis.

The presence and the composition of atherosclerotic plaques in the carotid artery are predictive of stroke and coronary events (Bos et al., 2021). The *degree of stenosis* is a commonly-used biomarker that assesses the severity of atherosclerosis of the carotid artery and is used in the decision-making for an endarterectomy procedure (Bonati et al., 2021). One way to measure the degree of stenosis is as defined by the European Carotid Artery Trial (ECST) (European Carotid Surgery Trialists' Collaborative Group, 1998) and requires the annotations of the diameter of the lumen and of the full vessel at the maximum narrowing. In this paper we explore whether these widely available clinically relevant diameters can be used to train segmentation models directly, without requiring additional voxelwise segmentations. The resulting segmentations enable the assessment of various imaging biomarkers, including the degree of stenosis, as well as measurements such as intima-media plaque thickness and total plaque volume.

This paper is a follow-up of Camarasa et al. (2022) presented at MIDL 2022. We extend the original concept of differentiable boundary

---

(a) Segmentation probability map   (b) Centroid of the inner class   (c) Inner and outer classes radii estimates   (d) Diameter estimates at maximum stenosis
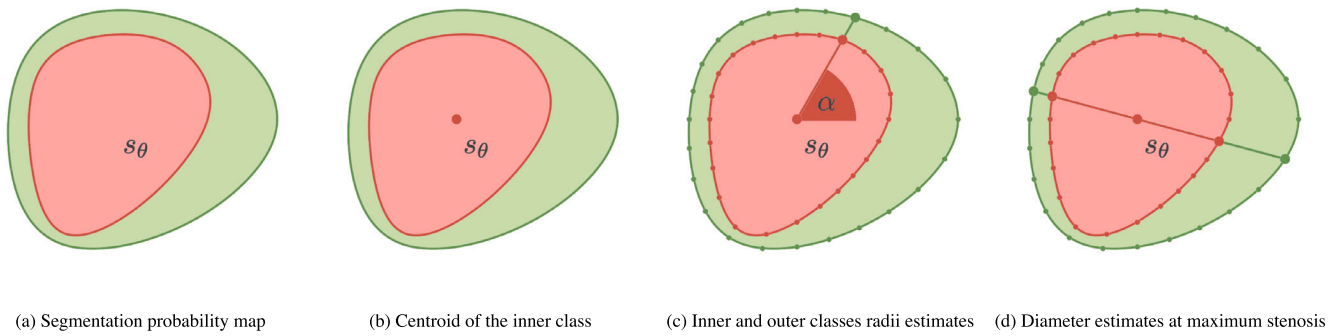
Fig. 1. Method pipeline. The inner class ($\mathcal{I}$ in red), the outer class ($\mathcal{O}$ in green) and the foreground classes (as the union of inner and outer classes $\overline{\mathcal{B}} = \mathcal{I} \cup \mathcal{O}$).

point extraction in a multi-class setting—which in the case of carotid arteries allows to segment both the artery lumen and wall, providing more relevant clinical information—and provide more extensive validation. Our contributions can be summarized as follows:

- we extend the differentiable boundary point extraction developed in Camarasa et al. (2022) to the multi-class segmentation of two nested star-shaped objects;
- we optimize this multi-class segmentation problem using diameter annotations extracted at the maximum narrowing;
- we comprehensively evaluate our method on carotid artery segmentation in two different MR image datasets of multi-center studies of patients who suffered a recent vascular event.

2. Related work

2.1. Weak supervision

Weakly supervised segmentation has received increased attention in recent years. Most methods use "standard" labels: bounding boxes (Papandreou et al., 2015; Rajchl et al., 2016; Kervadec et al., 2020), scribbles (Lin et al., 2016; Wang et al., 2018), dots (Bearman et al., 2016; Qu et al., 2019; Luo et al., 2021; Dorent et al., 2021) or image-level labels (image tags) (Pathak et al., 2015; Sahasrabudhe et al., 2020). Most of those methods are generic and can be applied to a variety of tasks without further modification.

When strong priors on the segmented object exist, some weakly supervised methods can be tailored to efficiently exploit this information and better compensate for the lack of full annotations. For instance, a volume prior can be enforced through a quadratic penalty on the size of the segmented object (Kervadec et al., 2019b) or the proportion of voxels belonging to the foreground (Bortsova et al., 2018). Sahasrabudhe et al. (2020) and Qu et al. (2019) exploit the structure of deep nuclei segmentation by either learning to predict the image zoom-levels (as a proxy-task), or building Voronoi cells from the dot annotations, respectively. Dorent et al. (2021) assumes that the object is continuous and "draws" a path between the extreme points annotations, by minimizing a geodesic distance. In a multi-instance object detection in pathology images, Yang et al. (2020) detects circles using center and radius information at train time, which can be seen as a segmentation with a strong circular prior.

Two existing methods could be tweaked for the aforementioned task of carotid artery segmentation when using only a diameter annotation. InExtremIS (Dorent et al., 2021) can use the diameter as a valid path without resorting to the complex geodesic distance calculation; but it would not exploit the diameter information. CircleNet (Yang et al., 2020) can be modified to predict a single circular segmentation, using the center and diameter information but ignoring the exact boundary of the object. As such, to the best of our knowledge, there exists no method in the current literature that could use the information on the diameters of a vessel lumen and the outer wall while segmenting the object precisely.

2.2. Carotid artery segmentation

Considering the segmentation of the carotid artery lumen in magnetic resonance (MR) images, Luo et al. (2019) exploited its hyper-intensity on TOF-MR images developing a variation of the level set method. Segmenting both the lumen and the outer wall, Arias Lorza et al. (2018), Arias-Lorza et al. (2016) used a geometrical prior of the artery relying on an optimal surface graph-cut algorithm. Both these methods required a prior estimate of the artery centerline.

The rise of deep learning methods moved research away from those semi-automatic algorithms towards more generalizable, fully-automatic, and end-to-end optimizable methods that could apply to many (if not all) segmentation tasks (Isensee et al., 2021). As an example, multiple versions of the well-established U-Net network (Ronneberger et al., 2015) trained with voxelwise annotations reported good performances of the joint segmentation of the lumen and the outer wall of the carotid artery: Wu et al. (2019) compared different U-Net variations on 2D T1 weighted MR images, Zhu et al. (2021) developed a cascading approach of residual U-Net for 3D MR and Camarasa et al. (2021) investigated performances and the uncertainty of the prediction of a deep learning method trained on a multi-sequence MR images dataset.

Combining a naïve-Bayesian method with a level-set method, Liu et al. (2006) segmented the different components of the plaque while van Engelen et al. (2015) preferred a support vector machine model. Zhang et al. (2019) compared four established machine learning models (random forest, artificial neural network, gradient boosting decision tree, and artificial neural network) applied to the pixel-wise classification of the plaque components.

In 2021, a carotid artery segmentation challenge was held based on MR imaging data of the Care II study (Zhao et al., 2017). The participants of the challenge were tasked to automatically segment the outer wall of the internal, external, and common carotid arteries on 3D black blood MR images of 24 patients. Alblas et al. (2021) won the challenge with a two-stage approach that aimed to locate the center-line and subsequently estimate the lumen and outer-wall contours.

To the best of our knowledge, little research focuses on using relevant weak labels to supervise carotid artery MR segmentation algorithms.

3. Method

In this section, we present a method to approximate the radii (Section 3.1), the diameters (Section 3.2), and the centroid (Section 3.3) of a star-shaped object. We introduce a weakly-supervised segmentation setting based on diameter annotations (Section 3.4) where we apply our methodology to the optimization of star shaped objects segmentation probability maps produced by a deep learning model (Section 3.5). For convenience, all mathematical notation is summarized in Appendix A.1, though each symbol is introduced in time in the manuscript.
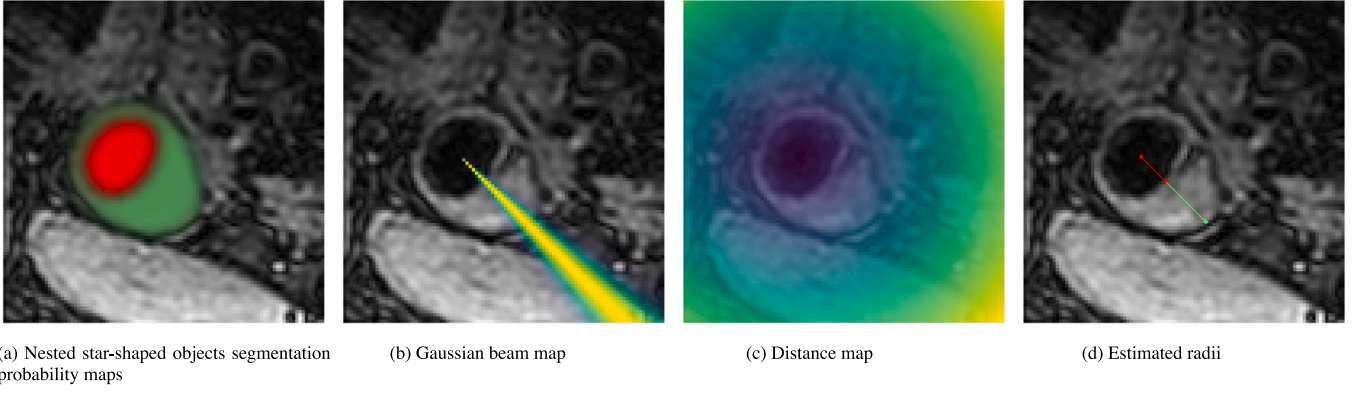
(a) Nested star-shaped objects segmentation probability maps

(b) Gaussian beam map

(c) Distance map

(d) Estimated radii

Fig. 2. β-parameterized estimates of the inner and foreground radii at a given angle ($\beta = 1$).

3.1. Estimating the radius of a star-shaped object

Definition 3.1.1 (*Star-Shaped Object*).

An object (S) of an imaging domain (Ω) is star-shaped if there exists a non-empty set of *root points* (J_S) such that any ray originating from a *root point* ($j \in J_S$) crosses the boundary (∂S) of the object exactly once. This translates mathematically, as follows:

$$\forall S \in \mathbb{S}, \quad \underbrace{\exists J_S \subset_{\neq \emptyset} S}_{\substack{\text{non-empty set} \\ \text{of root points}}} \quad :$$

$$j \in J_S \implies \forall \alpha \in [0, 2\pi[, \left| \underbrace{R_j(\alpha)}_{\text{ray}} \cap \overbrace{\partial S}^{\text{boundary}} \right| = 1 \tag{1}$$

where $\mathbb{S}$ is the set of star-shaped objects of Ω, and $R_j(\alpha) = \{j + a.(\cos(\alpha), \sin(\alpha)) : a \in \mathbb{R}^+\}$ the ray originated from j and of angle α.

Definition 3.1.2 (*Radius*). For a given orientation $\alpha \in [0, 2\pi[$ and a *root point* $j \in J_S$, we define the radius $\mathfrak{R}(\mathbf{1}_S, \alpha, j)$ of the star-shaped objects $S \in \mathbb{S}$ as the distance between the *root point* j and the intersection of the boundary of the star-shaped object ∂S and the ray of angle α originating from j:

$$\mathfrak{R}(\mathbf{1}_S, \alpha, j) = \text{dist}\left(j, R_j(\alpha) \cap \partial S\right) \tag{2}$$

where $\mathbf{1}_S$ is the indicator function of the star-shaped object S and dist the Euclidean distance between a point and a set. As S is star-shaped (see Definition 3.1.1), the set $R_j(\alpha) \cap \partial S$ contains a single element.

Definition 3.1.3 (*Gaussian Beam*). Considering the change of coordinates $\Phi_j : u \to (\Phi_j(u)_r, \Phi_j(u)_\gamma)$ that transforms a Cartesian coordinate u into a radial coordinate $\Phi_j(u)_r$ and an angular coordinate $\Phi_j(u)_\gamma$, we define a Gaussian beam:

$$f : \begin{cases} \Omega \times [0, 2\pi[\times J_S \to \mathbb{R}^+, \\ (u, \alpha, j) \mapsto g_\alpha(\Phi_j(u)_\gamma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{\Phi_j(u)_\gamma - \alpha}{2\sigma^2}\right). \end{cases} \tag{3}$$

where g_α is the probability density function of a Gaussian distribution centered on α with standard deviation σ.

Fig. 2(b) offers a visualization of the Gaussian beam displayed as a map overlaid on an MR image.

Proposition 3.1.1 (*β-Parameterized Radial Estimate*). *Given a star-shaped object $S \in \mathbb{S}$, a root point $j \in J_S$, an orientation $\alpha \in [0, 2\pi[$ and a parameter $\beta \in \mathbb{R}^+$, under the assumptions of:*

- *local circularity (on the angular interval $[\alpha - 3\sigma, \alpha + 3\sigma]$, S can be considered circular): $\forall \gamma \in [\alpha - 3\sigma, \alpha + 3\sigma] : \mathfrak{R}(\mathbf{1}_S, \gamma, j) \approx \mathfrak{R}(\mathbf{1}_S, \alpha, j)$;*

- *concentration of the beam (the vast majority of the distribution g_α is concentrated within the angular interval $[\alpha - 3\sigma, \alpha + 3\sigma]$): $\int_{\gamma \in \mathbb{R}} g_\alpha(\gamma)d\gamma \approx \int_{\alpha-3\sigma}^{\alpha+3\sigma} g_\alpha(\gamma)d\gamma \approx 1,$*

a function of the radius can be approximated by an integral:

$$\underbrace{\int_{u \in \Omega} f(u, \alpha, j)\mathbf{1}_S(u)\|u - j\|_2^\beta du}_{\mathscr{I}_\beta(\mathbf{1}_S, \alpha, j)} \approx \frac{\mathfrak{R}(\mathbf{1}_S, \alpha, j)^{\beta+2}}{\beta + 2}, \tag{4}$$

giving a β-parameterized radial estimate of the true radius:

$$\widehat{\mathfrak{R}}_\beta(\mathbf{1}_S, \alpha, j) = \left((\beta + 2)\mathscr{I}_\beta(\mathbf{1}_S, \alpha, j)\right)^{\frac{1}{\beta+2}}. \tag{5}$$

The proof of Proposition 3.1.1 can be found in Appendix A.3. The β-parameterized radial estimates of two nested star-shaped objects can then be obtained by integrating the voxelwise product of the Gaussian beam (f), the segmentation probability map ($\mathbf{1}_S$) and a distance map ($\|u - j\|^\beta$). This process is illustrated in Fig. 2.

Proposition 3.1.2 (*Radial Estimate*). *Considering a star-shaped object $S \in \mathbb{S}$, a root point $j \in J_S$, an orientation $\alpha \in [0, 2\pi[$, and a set of β-parameters $B \subset \mathbb{R}_+$, averaging the available β-parameterized radial estimates provides a radial estimate:*

$$\widehat{\mathfrak{R}}(\mathbf{1}_S, \alpha, j) = \mathbb{E}_{\beta \in B}\left(\widehat{\mathfrak{R}}_\beta(\mathbf{1}_S, \alpha, j)\right) \tag{6}$$

3.2. Estimating diameters

Definition 3.2.1 (*Star-Shaped Diameter*). Given a star-shaped object $S \in \mathbb{S}$, a root point $j \in J_S$, an orientation $\alpha \in [0, \pi[$, the diameter $\mathfrak{D}$ is defined as the sum of two opposite radii:

$$\mathfrak{D}(\mathbf{1}_S, \alpha, j) = \mathfrak{R}(\mathbf{1}_S, \alpha, j) + \mathfrak{R}(\mathbf{1}_S, \alpha + \pi, j). \tag{7}$$

Proposition 3.2.1 (*Diameter Estimate*). *Considering a star-shaped object $S \in \mathbb{S}$, a root point $j \in J_S$ and an orientation $\alpha \in [0, \pi[$, a diameter estimate $\widehat{\mathfrak{D}}$ can be defined as the sum of two opposite radial estimates:*

$$\widehat{\mathfrak{D}}(\mathbf{1}_S, \alpha, j) = \widehat{\mathfrak{R}}(\mathbf{1}_S, \alpha, j) + \widehat{\mathfrak{R}}(\mathbf{1}_S, \alpha + \pi, j). \tag{8}$$

Notice that these definitions of the diameter and of its estimate differ from Camarasa et al. (2022) as they force a certain *root point* j to belong to the diameter. This modification is motivated by the available annotations and the prior knowledge of our specific application presented in Section 3.4.

3.3. Centroid

Definition 3.3.1 (*Centroid*). We define the centroid $\mathfrak{c} : (\Omega \to [0, 1]) \to \Omega$ of a set $A \subset \Omega$ as the average of the coordinates of the set:

$$\mathfrak{c}\left(\mathbf{1}_A\right) = \frac{1}{\int_{u \in \Omega} \mathbf{1}_A(u)\mathrm{d}u} \int_{u \in \Omega} u.\mathbf{1}_A(u)\mathrm{d}u, \tag{9}$$

where $\mathbf{1}_A$ is the indicator function of the set A.

The more general notion of the centroid (not limited to star-shaped objects) was already investigated in the literature and shown successful in the weak-supervision of medical image segmentation (Kervadec et al., 2021). In our application, it will play an important role as *root point* of star-shaped objects (Section 3.4) and provide spatial information to our deep learning model (Section 3.5).

3.4. Application

In the following, we apply the previously presented theory to the segmentation of multiple registered sequences $x : \Omega \subset \mathbb{R}^2 \to \mathbb{R}^M$ in three non-overlapping subsets (the background subset ($\mathcal{B}$), the inner subset ($\mathcal{I}$) and the outer subset ($\mathcal{O}$)):

$$\mathcal{B} \cup \mathcal{I} \cup \mathcal{O} = \Omega$$

In our setting, we do not have access to the voxelwise ground-truth

$$y : \begin{cases} \Omega & \to & \Delta_2 \\ u & \mapsto & (\mathbf{1}_{\mathcal{B}}(u), \mathbf{1}_{\mathcal{I}}(u), \mathbf{1}_{\mathcal{O}}(u)) \end{cases}$$

but two pairs of annotated landmarks on the objects boundaries $(l'_{\overline{\mathcal{B}}}, l''_{\overline{\mathcal{B}}}) \in \partial \overline{\mathcal{B}}^2, (l'_{\mathcal{I}}, l''_{\mathcal{I}}) \in \partial \mathcal{I}^2$, indicating the diameters and therewith the degree of stenosis (European Carotid Surgery Trialists' Collaborative Group, 1998), (see Fig. 1).The annotated landmark pairs:

1. are *co-linear*: $\left| \left\langle l'_{\overline{\mathcal{B}}} - l''_{\overline{\mathcal{B}}}, l'_{\mathcal{I}} - l''_{\mathcal{I}} \right\rangle \right| = \left\| l'_{\overline{\mathcal{B}}} - l''_{\overline{\mathcal{B}}} \right\| . \| l'_{\mathcal{I}} - l''_{\mathcal{I}} \|$;
2. are *centered*: $\forall (l_1, l_2) \in \left\{ l'_{\overline{\mathcal{B}}}, l''_{\overline{\mathcal{B}}}, l'_{\mathcal{I}}, l''_{\mathcal{I}} \right\}, |\langle l_1 - \mathfrak{c}(\mathbf{1}_{\mathcal{I}}), l_2 - \mathfrak{c}(\mathbf{1}_{\mathcal{I}}) \rangle | = \| l_1 - \mathfrak{c}(\mathbf{1}_{\mathcal{I}}) \| . \| l_2 - \mathfrak{c}(\mathbf{1}_{\mathcal{I}}) \|$
3. maximize the *narrowing,* i.e. they minimize the ratio between the inner and outer diameter: $\left(l'_{\overline{\mathcal{B}}}, l''_{\overline{\mathcal{B}}}, l'_{\mathcal{I}}, l''_{\mathcal{I}} \right) \mapsto 1 - \frac{\| l'_{\mathcal{I}} - l''_{\mathcal{I}} \|_2}{\left\| l'_{\overline{\mathcal{B}}} - l''_{\overline{\mathcal{B}}} \right\|_2}$.

We also have the following prior knowledge about the subsets topology and nesting (see Fig. 1):

- the inner subset is a star-shaped object: $\mathcal{I} = y^{-1}([0, 1, 0]) \in \mathbb{S}$;
- the foreground subset (corresponding to the union of the inner and outer subsets) is a star-shaped object: $\overline{\mathcal{B}} = \mathcal{I} \cup \mathcal{O} = y^{-1}(\{[0, 1, 0]\}) \cup y^{-1}(\{[0, 0, 1]\}) \in \mathbb{S}$;
- the centroid of the inner subset is a *root point* of both the inner and the foreground subsets: $\mathfrak{c}(\mathcal{I}) \in J_{\overline{\mathcal{B}}} \cap J_{\mathcal{I}}$;

Our goal is to approximate the unavailable ground truth y with a deep learning model:

$$s_\theta : \begin{cases} \Omega & \to & \Delta_2 \\ u & \to & (\mathcal{B}_\theta(u), \mathcal{I}_\theta(u), \mathcal{O}_\theta(u)) \end{cases}$$

using the annotated landmarks and the available prior knowledge.

3.5. Optimization

As the centroid ($\mathfrak{c}$), the diameter estimate ($\hat{\mathfrak{d}}$) and the β-parameterized radial estimate ($\hat{\mathcal{R}}_\beta$) are differentiable, they can be used for the optimization of our deep learning model (s_θ) (see Fig. 1).

With a mean squared error loss, the predicted inner class $\mathcal{I}$ centroid can be learned using the inner class annotated centroid $\left(\frac{1}{2}(l'_{\mathcal{I}} + l''_{\mathcal{I}}) \right)$ and similarly the inner $\mathcal{I}$ and foreground $\overline{\mathcal{B}}$ classes diameters estimates using the annotated diameters ($\left\| l'_{\mathcal{I}} - l''_{\mathcal{I}} \right\|_2$ and $\left\| l'_{\overline{\mathcal{B}}} - l''_{\overline{\mathcal{B}}} \right\|_2$) as reference:

$$\mathcal{L}_{\mathfrak{c}}(\theta) = \left\| \log(\mathfrak{c}(\mathcal{I}_\theta)) - \log\left(\frac{l'_{\mathcal{I}} + l''_{\mathcal{I}}}{2} \right) \right\|_2^2, \tag{10}$$

$$\mathcal{L}_{\mathfrak{D}}(\theta) = \left\| \log\left(\hat{\mathfrak{d}}(\mathcal{I}_\theta, \alpha_l, \mathfrak{c}(\mathcal{I}_\theta)) \right) - \log\left(\left\| l'_{\mathcal{I}} - l''_{\mathcal{I}} \right\|_2 \right) \right\|_2^2$$

$$+ \left\| \log\left(\hat{\mathfrak{d}}\left(\overline{\mathcal{B}}_\theta, \alpha_l, \mathfrak{c}(\mathcal{I}_\theta) \right) - \hat{\mathfrak{D}}(\mathcal{I}_\theta, \alpha_l, \mathfrak{c}(\mathcal{I}_\theta)) \right) \right.$$

$$\left. - \log\left(\left\| l'_{\overline{\mathcal{B}}} - l''_{\overline{\mathcal{B}}} \right\|_2 - \left\| l'_{\mathcal{I}} - l''_{\mathcal{I}} \right\|_2 \right) \right\|_2^2. \tag{11}$$

where $\alpha_l = \cos^{-1}\left(\frac{\langle l''_{\mathcal{I}} - l'_{\mathcal{I}}, (1,0) \rangle}{\| l''_{\mathcal{I}} - l'_{\mathcal{I}} \|_2} \right)$ is the orientation of the annotations. Notice that contrary to Camarasa et al. (2022), logarithms are applied to stabilize the optimization limiting its range of values.

As we will show in the experiments section, additional regularization is desirable. Favoring the agreement of multiple β-parameterized radial estimates will reward *locally circular*, binarized ($S_\theta \approx \mathbf{1}_S$) star-shaped objects segmentation probability maps, and therefore efficiently achieve this goal.

This mathematically translates into a minimization of the variance of the ratio of the β-parameterized radial estimates and the radial estimates:

$$\mathcal{R}(\theta) = \sum_{S_\theta \in \{\mathcal{I}_\theta, \overline{\mathcal{B}}_\theta\}} \mathbb{E}_{\alpha \in \Gamma} \left(\text{Var}_{\beta \in B} \left(\frac{\hat{\mathcal{R}}_\beta(S_\theta, \alpha, \mathfrak{c}(\mathcal{I}_\theta))}{\hat{\mathcal{R}}(S_\theta, \alpha, \mathfrak{c}(\mathcal{I}_\theta))} \right) \right). \tag{12}$$

where $\Gamma \subset [0, 2\pi[$ is a discrete set of angles and $B \subset \mathbb{R}_+$ a discrete set of β-parameters.

The final combined, weakly-supervised loss is, therefore:

$$\mathcal{L}(\theta) = \lambda \mathcal{L}_{\mathfrak{c}}(\theta) + \mu \mathcal{L}_{\mathfrak{D}}(\theta) + \nu \mathcal{R}(\theta), \tag{13}$$

with λ, μ and ν hyper-parameters balancing the different components.

4. Experiments

In this section, we apply the weakly-supervised method presented in Section 3 to the segmentation of the carotid artery lumen (inner subset $\mathcal{I}$) and outer wall (outer subset $\mathcal{O}$) on two datasets of MR images presented in Section 4.1. Section 4.3 presents our comparison to the state of the art and an ablation study. Finally, the implementation details can be found in Section 4.4.

4.1. Datasets

Care II study. We used a subset of the data from the Care II study (Zhao et al., 2017), a multi-center study which enrolled patients had a recent ischaemic stroke or transient ischaemic attack. The subset of 24 MRI scans of 24 patients was made available in the context of a carotid artery segmentation challenge.[1] The scans are 3D Motion Sensitized Driven Equilibrium prepared Rapid Gradient Echo, 3D-MERGE. Full lumen segmentations of either left or right internal and common carotid artery are available in on average 12.1% of the slices of a scan. This resulted in 2151 annotated 2D slices over the whole dataset.

PARISK study. We used carotid artery MR images acquired within the multi-center, multi-scanner, multi-sequence PARISK study (Truijman et al., 2014), a large prospective study to improve risk stratification in patients with mild to moderate carotid artery stenosis who recently had a transient ischaemic attack, an amaurosis fugax, or a minor stroke. We used the images of 191 enrolled patients of the four study centers: Amsterdam Medical Center (AMC), the Maastricht University Medical Center (MUMC), the University Medical Center of Utrecht (UMCU), and Erasmus MC (EMC), all in the Netherlands. AMC, MUMC and UMCU performed the MR imaging with a 3.0-Tesla MR scanner with an eight-channel phased-array coil (Shanghai Chenguang Medical Technologies Co., Shanghai) and EMC used a dedicated four-channel carotid phased-array coil with an angulated setup (Machnet B.V., Roden, the Netherlands). UMCU and MUMC acquired the imaging data of all the patients with an Achieva TX scanner (Phillips Healthcare, Best, Netherlands), AMC center acquired the scans of 11 of its patients

---

[1] https://vessel-wall-segmentation.grand-challenge.org/

with an Ingenia scanner (Phillips Healthcare, Best, Netherlands), and 2 with an Intera scanner (Phillips Healthcare, Best, Netherlands) and EMC acquired the imaging data of all patients with a Discovery MR 750 system (GE Healthcare Milwaukee, MI, USA). Each enrolled patient underwent 5 MR sequences. For the AMC, UMCU, and MUMC centers the imaged sequences were: T1 weighted quadruple inversion recovery (IR) turbo spin echo (SE) (pre- (a.) and post-contrast (b.)), T2 weighted turbo SE (c.), IR turbo field echo (FE) (d.), and time of flight fast FE (e.). For the EMC center patients, the imaged sequences were T1 weighted double IR fast SE (pre- (a.) and post-contrast (b.)), T2 weighted turbo SE (c.), spoiled gradient echo (d.), and fast spoiled gradient echo (e.). The letters in between parenthesis (e.g. (a.)) indicates the matching of similar sequences across centers. More details on the image sequences can be found in Truijman et al. (2014). MR sequences were semi-automatically, first affinely and then elastically registered to the T1 weighted pre-contrast sequence. The vessel lumen and outer wall were annotated manually slice-wise, by trained observers with approximately 3 years of experience, in the T1 weighted pre-contrast sequence. Registration and annotation were achieved with VesselMass software.[2] The observer annotated the common and internal carotid arteries (either left or right) where the clinical symptoms [of atherosclerosis] occurred.

4.2. Preprocessing and diameter annotation

The following subsection presents the processing common to Care II and PARISK datasets.

Normalization. Input images were normalized at a slice level to match a mean intensity value of zero and standard deviation of one using the official implementation of MONAI (Cardoso et al., 2022).

Padding. Images were padded with zeros to a common dimension for their number of voxels along a dimension to match the smallest possible multiple of $2^5 = 32$voxels. This results in padded images of dimension 576×576 px^2 for PARISK and 768×160 px^2 for Care II.

Cropping. At training and evaluation time all 2D slices are cut in half to have only one carotid artery per sub-image.

Diameter annotations. We simulated the lumen and the full vessel diameter annotations $((l'_S, l''_S))_{I,\overline{B}}$ slice-wise, with the following methodology: extraction of the lumen centroid, resampling of the boundary point in $N = 24$ points (to ensure the *co-linearity* of the lumen and the full vessel diameters) and selection of the diameters in the direction of maximum *narrowing* (as defined in the method Section 3.4). Examples of the resulting simulated annotations can be found in Fig. 5 of the Appendix.

4.3. Baseline and ablation study

We evaluated our method and the different baselines in a 2D setting with the lumen and the full-vessel diameters annotated per slice on both the PARISK dataset and the Care II dataset. We compare our method to InExtremIS (Dorent et al., 2021) and CircleNet (Yang et al., 2020), as they are the most related weakly supervised methods and most relevant to our problem. To perform a fair comparison across all methods and assess the influence of the annotation-based losses, the additional CRF loss (Tang et al., 2018) used by InExtremeIS is not included as all methods could benefit from it. To determine the non-supervised area by InExtremIS, we consider a circle that has the centroid of the lumen annotations ($\frac{1}{2}(l'_I + l''_I)$) as a center and a diameter of $2\delta + \|l'_{\overline{B}} - l''_{\overline{B}}\|_2 - \delta$ (PARISK: $\delta = 7$ voxels and Care II study: $\delta = 4$ voxels) corresponds to the 95th percentile of the slice-wise difference between the maximum full vessel radius and the radius at the maximum narrowing (computed

over the whole dataset). For CircleNet, the predicted heatmap and the radius map spatial dimension match the spatial dimension of the input images rather than being down-sampled as originally proposed by Yang et al. (2020). This slight modification enables us to use the same network architecture for all methods and therefore make a more fair comparison. Additionally, we compared to a fully-supervised U-Net as an upper-bound, and perform an ablation study on the different components of the loss in Eq. (13).

4.4. Implementation details

All methods are built on top of the base U-Net architecture (Ronneberger et al., 2015) using the official MONAI 0.7 implementation as starting point (Cardoso et al., 2022). We trained the model with the same ADAM optimizer (learning rate 10^{-4}, $\beta_1 = 0.9$, $\beta_2 = 0.99$) for 1000 epochs (for the Care II dataset) and 2000 epochs (for the PARISK dataset). We increased the number of epochs for PARISK as after 1000 epochs all methods did not reach full convergence. Our loss components (Eqs. (10), (11) and (12)) do not require any modification of the network architecture or training regime and are implemented as direct losses.

We use $|\Gamma| = 24$ radii equally spread between $[0, 2\pi[$ and $B = \{0,1\}$ as β-parameters. λ, μ, ν, and σ have been empirically set to 100, 100, 10, and 0.15 such that the contributions of the different loss terms are approximately equal, based on the average amplitude of the components of the loss at training time. As the models are trained to segment a single vessel, for all methods the final segmentation is chosen as the largest connected component of the network binarized foreground output (using the argmax operation).

A more extensive explanation of our implementation details can be found in our publically available repository of code at https://gitlab.com/radiology/aim/carotid-artery-image-analysis/nested-star-shaped-objects

4.5. Metrics and evaluation

All methods were trained and evaluated using 4-fold cross-validation (the folds are determined at a patient level): 2 folds for training, 1 fold for validation, and 1 fold for testing. Evaluation is performed per-slice with Dice score (DSC), Hausdorff distance (HD), and absolute diameter error in the direction of the vessel's maximum stenosis (ADE). We report the per-patient averages in the testing sets and perform a two-sided Wilcoxon signed-rank test, with a level of significance of 0.05 to determine if a baseline (or ablation) is significantly different from our proposed method.

5. Results

In this section, we present the results of the ablation study (Section 5.1) and of the comparison to the baseline methods (Section 5.2). For both the ablation study and the comparison to the baselines, we can find the quantitative results in Tables 1 and 2, the qualitative results in Fig. 3 and the influence of the post-processing in Fig. 4.

5.1. Ablation study

The ablation study shows, on both Care II and PARISK datasets, that each component of the loss is important. Removing the regularizer, ($\mathfrak{c} + \mathfrak{d}$) still allows the network to locate the lumen and to start to retrieve the shape of the vessel but it under-segments parts of both the lumen and outer wall, while supervising using only the centroid loss ($\mathfrak{c}$) gives bad segmentation performances (Fig. 3).

Fig. 4 demonstrates that keeping the largest component as post-processing benefits less to our method ($\mathfrak{c} + \mathfrak{d} + \mathcal{R}$) than our method without regularizer ($\mathfrak{c} + \mathfrak{d}$).

---

[2] https://medisimaging.com/apps/vesselmass-re/

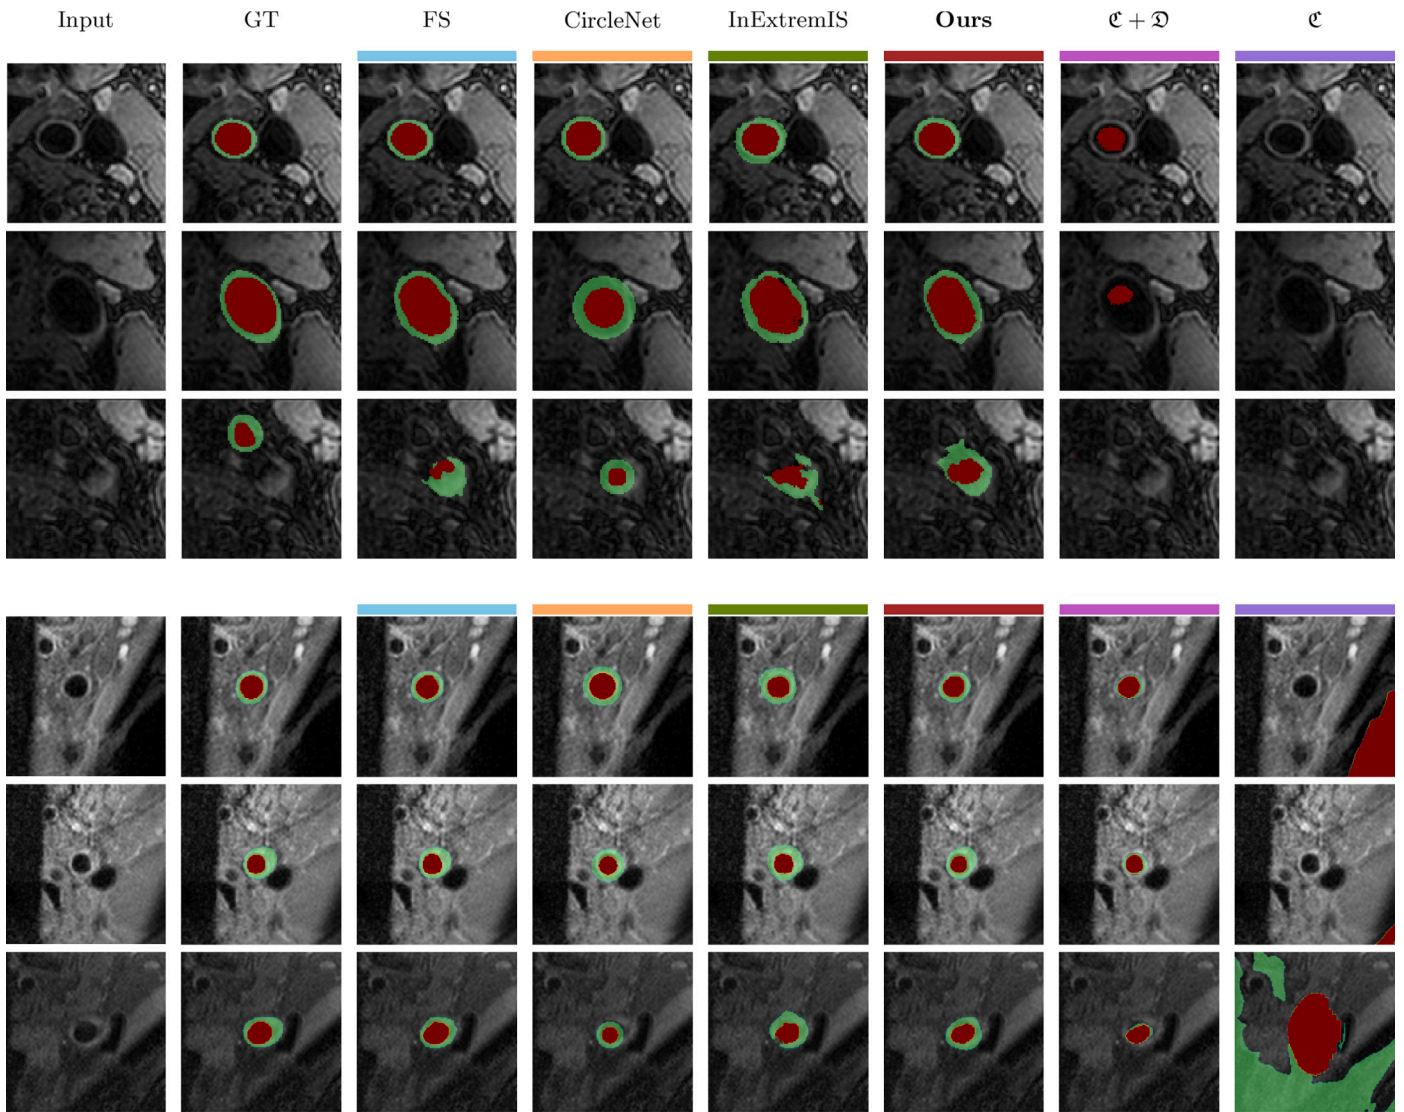| Input | GT | FS | CircleNet | InExtremIS | **Ours** | $\mathfrak{C} + \mathfrak{D}$ | $\mathfrak{C}$ |

Fig. 3. Lumen and outer wall segmentations for the different methods and different subjects. (top 3 rows: carotid artery challenge, bottom 3 rows: PARISK, GT: ground truth, FS: full supervision).

Table 1
Distribution of the metrics computed over the test set of Care II study dataset for the different methods. The reported values correspond to the median and the interquartile range (in between brackets). The results reported in bold the best performing method (apart from the full supervision) and the results denoted by a ∗ differ statistically from our proposed method using a Wilcoxon signed-rank test with a significance level of 0.05. (DSC: Dice Score, HD: Hausdorff Distance, ADE: Absolute Diameter Error).

	DSC ↑		HD (mm) ↓		ADE (mm) ↓	
	Lumen ($\mathcal{I}$)	Full vessel ($\overline{\mathcal{B}}$)	Lumen ($\mathcal{I}$)	Full vessel ($\overline{\mathcal{B}}$)	Lumen ($\mathcal{I}$)	Full vessel ($\overline{\mathcal{B}}$)
Full Supervision	0.85 *[0.82–0.90]*	0.85 *[0.82–0.91]*	2.9 *[1.3–7.4]*	2.4 *[1.3–4.5]*	0.76 *[0.59–0.87]*	0.82 *[0.67–1.0]*
Circlenet (Yang et al., 2020)	0.82∗ *[0.76–0.85]*	0.83 *[0.78–0.87]*	**1.6** *[1.3–2.7]*	**1.9** *[1.3–3.0]*	**0.71**∗ *[0.53–0.87]*	**0.84** *[0.67–1.0]*
InExtremIS (Dorent et al., 2021)	**0.85** *[0.78–0.89]*	0.77∗ *[0.66–0.81]*	2.3∗ *[1.5–3.8]*	3.1∗ *[2.2–4.6]*	0.68∗ *[0.57–0.85]*	1.6∗ *[1.3–2.2]*
$\mathfrak{C} + \mathfrak{D} + \mathcal{R}$ (Ours)	**0.85** *[0.80–0.88]*	**0.84** *[0.78–0.87]*	1.9 *[1.2–2.5]*	2.0 *[1.7–2.9]*	0.86 *[0.72–1.0]*	0.94 *[0.75–1.1]*
$\mathfrak{C} + \mathfrak{D}$	0.71∗ *[0.48–0.81]*	0.56∗ *[0.33–0.67]*	4.6∗ *[2.1–11.]*	5.8∗ *[3.0–12.]*	1.7∗ *[1.3–3.8]*	3.9∗ *[2.5–6.1]*
$\mathfrak{C}$	0.0036∗ *[0.00073–0.019]*	0.0013∗ *[0.00024–0.0092]*	1.3e+02∗ *[1.2e+02–1.4e+02]*	1.0e+02∗ *[97.–1.0e+02]*	17.∗ *[5.9–20.]*	1.0e+02∗ *[71.–1.0e+02]*

Table 2
Distribution of the metrics computed over the test set of PARISK dataset for the different methods. The reported values correspond to the median and the interquartile range (in between brackets). The results reported in bold the best performing method (apart from the full supervision) and the results denoted by a * differ statistically from our proposed method using a Wilcoxon signed-rank test with a significance level of 0.05. (DSC: Dice Score, HD: Hausdorff Distance, ADE: Absolute Diameter Error).

|	DSC ↑		HD (mm) ↓		ADE (mm) ↓	
	Lumen ($\mathcal{I}$)	Full vessel ($\overline{\mathcal{B}}$)	Lumen ($\mathcal{I}$)	Full vessel ($\overline{\mathcal{B}}$)	Lumen ($\mathcal{I}$)	Full vessel ($\overline{\mathcal{B}}$)
Full Supervision	0.86 *[0.74–0.93]*	0.83 *[0.72–0.91]*	2.1 *[0.78–5.2]*	2.7 *[1.3–5.2]*	0.71 *[0.50–1.0]*	1.2 *[0.78–1.9]*
Cirlenet (Yang et al., 2020)	0.64* *[0.51–0.77]*	0.71* *[0.57–0.82]*	3.5 *[1.7–6.3]*	3.7* *[2.1–6.6]*	**0.91*** *[0.72–1.2]*	**1.4*** *[1.0–1.9]*
InExtremIS (Dorent et al., 2021)	**0.83*** *[0.69–0.89]*	0.70* *[0.60–0.78]*	3.1 *[1.4–6.4]*	4.5* *[2.8–6.8]*	0.93* *[0.65–1.5]*	2.7* *[2.0–4.0]*
$\mathfrak{C}+\mathfrak{D}+\mathcal{R}$ (Ours)	0.77 *[0.55–0.85]*	**0.76** *[0.62–0.83]*	**2.2** *[1.3–6.2]*	**3.2** *[2.3–5.5]*	1.1 *[0.72–1.6]*	1.4 *[1.1–2.0]*
$\mathfrak{C}+\mathfrak{D}$	0.61* *[0.17–0.81]*	0.42* *[0.11–0.61]*	15.* *[1.3–1.2e+02]*	7.0* *[3.2–1.1e+02]*	1.5* *[0.93–2.9]*	4.3* *[2.3–7.0]*
$\mathfrak{C}$	$4.6e-08$* *[3.3e-08–1.2e-07]*	$1.0e-10$* *[9.4e-11–1.1e-10]*	2.0e+02* *[1.7e+02–2.0e+02]*	1.1e+02* *[1.0e+02–1.2e+02]*	23.* *[12.–34.]*	1.5e+02* *[1.2e+02–1.6e+02]*
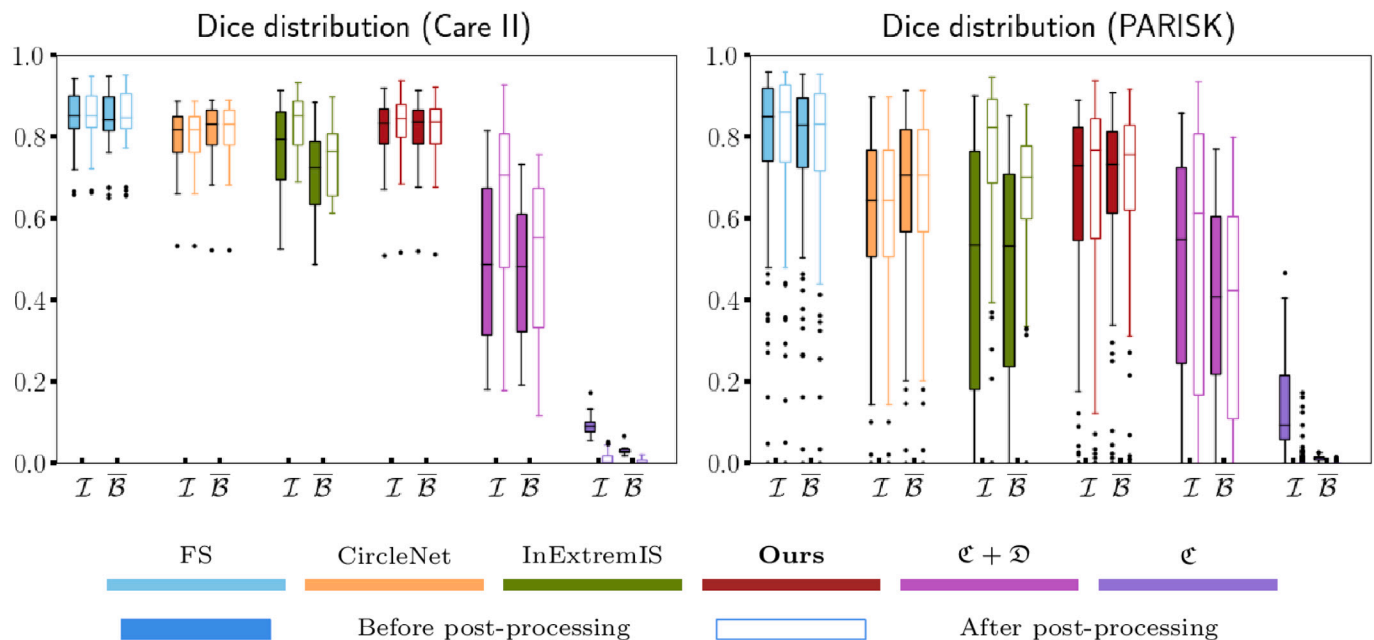
Fig. 4. Boxplots showing the Dice Score (averaged per patient) before and after post-processing for both the lumen ($\mathcal{I}$) and the full vessel ($\overline{\mathcal{B}}$). Each box goes from Q1 to Q3 while displaying the median. The whiskers extend from the box by 1.5 times the interquartile range.

5.2. Baselines

As could be expected, the full-supervision outperforms weakly-supervised methods in most metrics (Tables 1 and 2) and is the upper bound of our comparison. However, the best weakly-supervised methods come close to the full supervision result, most notably in terms of Dice for the Care II dataset (Table 1).

The third row of Fig. 3 shows one of the main challenges of both our proposed method and the considered baselines: segmenting the correct vessel when multiple are present in the image. In this example, all methods consistently segment the external instead of the internal carotid artery.

The proposed method shows, across datasets and metrics, significantly better segmentation of the full vessel than ($\overline{\mathcal{B}}$) InExtremIS (Tables 1 and 2). This translates qualitatively (Fig. 3) into a better understanding of the shape of the outer wall. In terms of lumen seg-

mentation, InExtremIS (after post-processing) seems to outperform our proposed method however not on every metric and dataset (Tables 1 and 2).

Our method segments the lumen of the carotid artery better (on both datasets) than CircleNet (Tables 1 and 2). In terms of full vessel segmentation, both methods seem to perform similarly (Tables 1 and 2). However, the qualitative results (Fig. 3) highlight the difficulties of CircleNet to capture more complex shapes such as elongated vessels close to the carotid artery bifurcation as it can only predict perfect circles.

It should be noted that postprocessing in the form of keeping only the largest connected component is crucial for InExtremIS, which has poor results without (Fig. 4). In contrast, this post-processing gives only a modest improvement for both the proposed method and the full supervision. CircleNet, by design outputs a single component and is unaffected by this step.
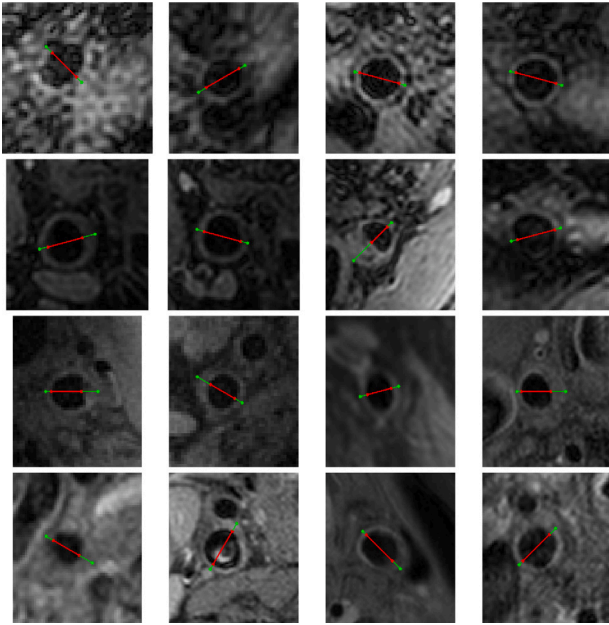
Fig. 5. Simulated diameters annotations of the lumen (I) and the full vessel ($\overline{B}$). (top 2 rows: Care II dataset, bottom 2 rows: PARISK dataset).
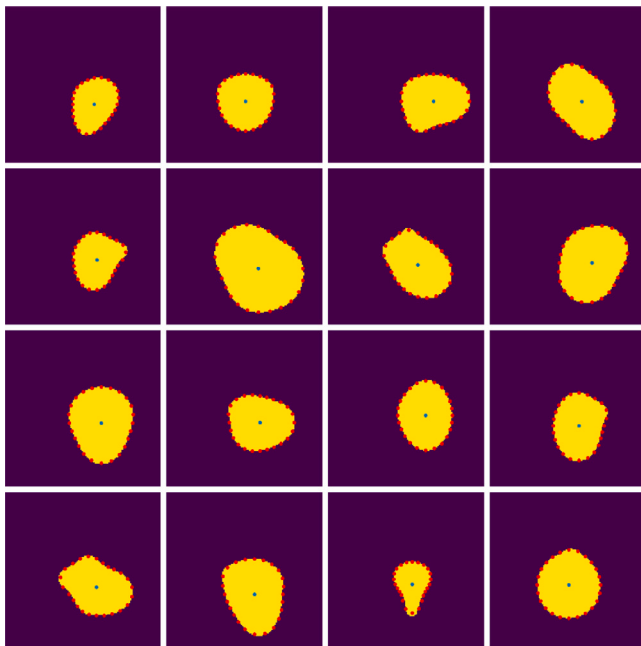

Fig. 6. Application of the differentiable boundary point extraction applied to randomly generated star-shaped objects (green: generated star-shape object, red: extracted boundary points).

6. Discussion

A simple modification of our previously proposed method (Camarasa et al., 2022) transformed the segmentation of a single star-shaped object using maximum diameter into the multi-class segmentation of nested star-shaped objects using diameter annotations in the orientation of maximum vessel narrowing.

Our ablation study shows the importance of the regularizer which has two main effects. First and foremost, it pushes the deep learning

model to predict a single, binarized, and locally circular object, making use of the prior knowledge at hand. Second, it increases the number of supervised voxels at each backward pass—the Gaussian beams used for the regularizer cover the entire image while the ones of the loss over the diameters concentrate in a shallow area in the direction of the vessel's maximum narrowing.

Our method demonstrates strong performances both quantitatively and qualitatively. It can capture more complex shapes than CircleNet (Yang et al., 2020) while being more stable and having a better understanding of the vessel structure than InExtremIS (Dorent et al., 2021).

A limitation of our study is that we used simulated diameter annotations. A subset of these simulated annotations was checked and approved by a medical doctor. Still, annotations made in clinical practice might not always be precisely at the maximum narrowing of the vessel, which could slightly reduce the performance of the proposed approach. However, we expect methods relying on the exact supervision of voxels (such as InExtremIS Dorent et al., 2021) to be even more sensitive to this type of uncertainty in the annotations.

In our setting, all methods have access to slice-wise diameters. In a more realistic setup, the two diameters, necessary to measure the degree of stenosis (European Carotid Surgery Trialists' Collaborative Group, 1998), would be available on only one slice of the whole volume.

Depending on the modality and available software, clinical experts could prefer to annotate diameters in 3D instead of 2D. Although not applicable out of the box to 3D diameters, the presented method can easily be adapted to this type of annotation. This could be achieved by changing the β-parameterized radial estimate to the spherical coordinates system instead of the polar coordinates system (Appendix A.3), and defining the Gaussian beam as a function of the azimuth and the inclination instead of the angular coordinate (Definition 3.1.3).

In this paper, we compared our weakly-supervised method to full supervision. The two approaches can also be combined as they do not require a modification of the network architecture. Similarly to E.L. Jurdi et al. (2021), Karimi and Salcudean (2019), Kervadec et al. (2019a), our proposed loss (Eq. (13)) could complement classic segmentation losses which usually optimize a voxel-wise classification. This way, result of fully supervised segmentation may be improved for (nested) star-shaped object(s) and especially in the case that the training data is limited. In a semi-supervised fashion, our proposed regularizer (Eq. (12)) could supervise unlabeled data as it can be computed without annotations.

Our proposed approach bridges segmentation probability map and boundary Cartesian coordinates in a differentiable manner. In our case, we use this bridge to derive simple objects: diameters. However, this opens the door to more advanced objects lying in the Cartesian space such as parametric curves (Burdin et al., 1996), Fourier-based shape descriptors (El-ghazal et al., 2012) etc. Those more advanced objects could be used to regularize the optimization of segmentation tasks based on the prior knowledge at hand.

We evaluate our approach in this paper on fairly simple shapes, as many vessel cross-sections are elliptical or almost circular, it could be used to model and supervise more complex shapes as well (Appendix A.4, Fig. 6), such as tumors (Menze et al., 2015). This would be of clinical relevance as the well-established RECIST criterion assesses the progression of tumors based on the evolution of their longest diameter (Schwartz et al., 2016).

7. Conclusion

We have introduced a fully differentiable approach to locate the centroid and boundary points of a star-shaped object from a segmentation probability map. The method was then successfully applied to train a segmentation neural network to segment nested star-shaped objects, supervised by diameter annotations, and regularized exploiting the

Table 3
Mathematical notation.

Symbol	Meaning
$\Omega \subset \mathbb{R}^2$	Two-dimensional image space;
$A \subset_{\neq\emptyset} B$	A is a non empty set of B;
$S \in \mathbb{S}$	A star-shaped object;
$\partial S \subset S$	The boundary of a star-shaped object;
$J_S \subset_{\neq\emptyset} \mathbb{S}$	The ensemble of root points of a star-shaped object S;
$R_j : [0, 2\pi[\to \mathbb{P}(\Omega)$	The ray originating from the root point j for a given angle;
$\text{dist} : \Omega \times \mathbb{P}(\Omega) \to \mathbb{R}^+$	The Euclidean distance between a point and a set;
$u \in \Omega$	Coordinates in the Cartesian image space;
$(r, \gamma) \in \Lambda_j = \mathbb{R}^+ \times [0, 2\pi[$	Polar coordinates, centered around j;
$\Phi_j : \Omega \to \Lambda_j$	Function going from Cartesian to polar coordinates;
$\mathbf{1}_A : \Omega \to \{0, 1\}$	The indicator function of a subset A of Ω;
$\mathfrak{C} : (\Omega \to [0, 1]) \to \Omega$	Function computing the centroid of a probability map;
$\alpha \in \Gamma \subset [0, 2\pi[$	Some angle out of a discrete set;
$\beta \in B \subset \mathbb{R}^+$	Some exponent out of a discrete set;
$\mathfrak{R} : (\Omega \to \{0, 1\}) \times [0, 2\pi[\times \Omega \to \mathbb{R}^+$	Exact radius of a star-shaped object for a angle and root point;
$\widehat{\mathfrak{R}}_\beta : (\Omega \to [0, 1]) \times [0, 2\pi[\times \Omega \to \mathbb{R}^+$	β-parameterized radial estimate of a star-shaped object for a given angle and a given root point;
$\widehat{\mathfrak{R}} : (\Omega \to [0, 1]) \times [0, 2\pi[\times \Omega \to \mathbb{R}^+$	Radial estimate of a star-shaped object for a given angle and a given root point;
$\widehat{\mathfrak{D}} : (\Omega \to [0, 1]) \times [0, \pi[\times \Omega \to \mathbb{R}^+$	Diameter estimate of a star-shaped object for a given angle and a given root point;
$x : \Omega \to \mathbb{R}^M$	Multiple registered images (corresponding to M MR sequences);
$\mathcal{B} \cup \mathcal{I} \cup \mathcal{O} = \Omega$	Set of classes to segment ($\mathcal{B}$: background class, $\mathcal{I}$ inner class, $\mathcal{O}$ outer class);
$\Delta_n \subset [0, 1]^{n+1}$	The simplex of dimension n;
$s_\theta : \Omega \to \Delta_2$	The softmax probabilities;
$y : \Omega \to \Delta_2$	The unavailable ground truth;
$\left((l'_I, l''_I), (l'_{\overline{B}}, l''_{\overline{B}})\right) \in \partial \mathcal{I}^2 \times \partial \overline{\mathcal{B}}^2$	Diameter annotations at the maximum narrowing, for the inner class ($\mathcal{I}$) and the foreground class ($\overline{\mathcal{B}}$);
$\alpha_l = \cos^{-1}\left(\frac{\langle l''_I - l'_I, (1,0) \rangle}{\|l''_I - l'_I\|_2} \right)$	The orientation of the annotations;
$f : \Omega \times [0, 2\pi[\times J_S \to \mathbb{R}^+$	The Gaussian beam;
$\sigma \in \mathbb{R}^+$	Hyper-parameter controlling the spread of the Gaussian beam;
$g_\alpha : \mathbb{R} \to \mathbb{R}^+$	A probability density function of a Gaussian distribution centered on α and standard deviation σ;
$(\lambda, \mu, \nu) \in \mathbb{R}^3$	Hyper-parameters balancing the combined loss

available shape-prior knowledge. This provides a mathematically sound way to re-use existing clinically relevant annotations. We validated the method on two datasets of MR images of carotid arteries—using only diameter annotations at training time and segmenting both the lumen and outer wall—and showed segmentation performance approaching that of full voxelwise supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All the code necessary for our analysis is made accessible. The analysis has been conducted on two datasets one public dataset and one private dataset.

Acknowledgments

Appendix. Supplementary material

A.1. Mathematical notation

In this article, we denote a function (f) as a mapping from an input set (A) to an output set (B) which associates to an element of the input set ($x \in A$) its corresponding image ($f(x) \in B$). This translates mathematically to:

$$f \begin{cases} A \to B \\ x \mapsto f(x) \end{cases}$$

The relevant mathematical notations can be found in Table 3.

A.2. Simulated annotations

See Fig. 5.

A.3. Proof of Proposition 3.1.1

Let us derive Eq. (4)

$$\mathcal{I}_\beta(\mathbf{1}_S, \alpha, j) = \int_{u \in \Omega} f(u, \alpha, j) \mathbf{1}_S(u) \|u - j\|_2^\beta \mathrm{d}u,$$

we apply the change of variable $(r, \gamma) = \Phi_j(u)$:

$$\mathcal{I}_\beta(\mathbf{1}_S, \alpha, j) = \int \int_{(r, \gamma) \in \Lambda_j} f(\Phi_j^{-1}(r, \gamma), \gamma, j) \mathbf{1}_S(\Phi_j^{-1}(r, \gamma)) r^\beta r \mathrm{d}r \mathrm{d}\gamma$$

$$= \int\int_{(r,\gamma)\in\Lambda_j} g_\alpha(\gamma)\mathbf{1}_S(\Phi_j^{-1}(r,\gamma))r^{\beta+1}\mathrm{d}r\mathrm{d}\gamma$$

$$= \int_{\gamma\in[0,2\pi]}\int_{r\in\mathbb{R}^+} g_\alpha(\gamma)\mathbf{1}_{[0,\Re(\mathbf{1}_S,\gamma,j)]}(r)r^{\beta+1}\mathrm{d}r\mathrm{d}\gamma$$

$$= \int_{\gamma\in[0,2\pi]}\int_{r=0}^{\Re(\mathbf{1}_S,\gamma,j)} g_\alpha(\gamma)r^{\beta+1}\mathrm{d}r\mathrm{d}\gamma$$

$$= \frac{1}{\beta+2}\int_{\gamma\in[0,2\pi]} g_\alpha(\gamma)\Re(\mathbf{1}_S,\gamma,j)^{\beta+2}\mathrm{d}\gamma.$$

under the assumptions of *local circularity* and *concentration of the beam*, we show that:

$$\mathcal{I}_\beta(\mathbf{1}_S,\alpha,j) \approx \frac{1}{\beta+2}\int_{\gamma\in[\alpha-3\sigma,\alpha+3\sigma]} g_\alpha(\gamma)\Re(\mathbf{1}_S,\alpha,j)^{\beta+2}\mathrm{d}\gamma$$

$$\approx \frac{\Re(\mathbf{1}_S,\alpha,j)^{\beta+2}}{\beta+2}.$$

A.4. Synthetic examples

See Fig. 6.

References

Alblas, D., Brune, C., Wolterink, J., 2021. Deep Learning-Based Carotid Artery Vessel Wall Segmentation in Black-Blood MRI Using Anatomical Priors. WorkingPaper, arXiv.org, SPIE Medical Imaging 2022.

Arias-Lorza, A.M., Petersen, J., van Engelen, A., Selwaness, M., van der Lugt, A., Niessen, W.J., de Bruijne, M., 2016. Carotid artery wall segmentation in multispectral MRI by coupled optimal surface graph cuts. IEEE Trans. Med. Imaging 35 (3), 901–911. http://dx.doi.org/10.1109/TMI.2015.2501751.

Arias Lorza, A.M., Van Engelen, A., Petersen, J., Van Der Lugt, A., de Bruijne, M., 2018. Maximization of regional probabilities using Optimal Surface Graphs: Application to carotid artery segmentation in MRI. Med. phys. 45 (3), 1159–1169.

Bearman, A., Russakovsky, O., Ferrari, V., Fei-Fei, L., 2016. What's the point: Semantic segmentation with point supervision. In: European Conference on Computer Vision. Springer, pp. 549–565.

Bonati, L.H., Kakkos, S., Berkefeld, J., de Borst, G.J., Bulbulia, R., Halliday, A., van Herzeele, I., Koncar, I., McCabe, D.J., Lal, A., et al., 2021. European Stroke Organisation guideline on endarterectomy and stenting for carotid artery stenosis. Euro. Stroke J. 6 (2), I–XLVII.

Bortsova, G., Dubost, F., Ørting, S., Katramados, I., Hogeweg, L., Thomsen, L., Wille, M., de Bruijne, M., 2018. Deep learning from label proportions for emphysema quantification. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2018. Springer International Publishing, Cham, ISBN: 978-3-030-00934-2, pp. 768–776.

Bos, D., Arshi, B., van den Bouwhuijsen, Q.J., Ikram, M.K., Selwaness, M., Vernooij, M.W., Kavousi, M., van der Lugt, A., 2021. Atherosclerotic carotid plaque composition and incident stroke and coronary events. J. Am. Coll. Cardiol. 77 (11), 1426–1435. http://dx.doi.org/10.1016/j.jacc.2021.01.038.

Burdin, V., Roux, C., Lefevre, C., Stindel, E., 1996. Modeling and analysis of 3-D elongated shapes with applications to long bone morphometry. IEEE Trans. Med. Imaging 15 (1), 79–91. http://dx.doi.org/10.1109/42.481443.

Camarasa, R., Bos, D., Hendrikse, J., Nederkoorn, P., Kooi, M.E., van der Lugt, A., de Bruijne, M., 2021. A quantitative comparison of epistemic uncertainty maps applied to multi-class segmentation. Mach. Learn. Biomed. Imaging (ISSN: 2766-905X) 1, URL https://melba-journal.org/papers/2021:013.html.

Camarasa, R., Kervadec, H., Bos, D., de Bruijne, M., 2022. Differentiable boundary point extraction for weakly supervised star-shaped object segmentation. In: Medical Imaging with Deep Learning. URL https://openreview.net/forum?id=whpBn0oadz.

Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al., 2022. MONAI: An open-source framework for deep learning in healthcare. arXiv preprint arXiv:2211.02701.

Dorent, R., Joutard, S., Shapey, J., Kujawa, A., Modat, M., Ourselin, S., Vercauteren, T., 2021. Inter extreme points geodesics for end-to-end weakly supervised image segmentation. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2021. Springer International Publishing, Cham, ISBN: 978-3-030-87196-3, pp. 615–624.

El-ghazal, A., Basir, O., Belkasim, S., 2012. Invariant curvature-based Fourier shape descriptors. J. Vis. Commun. Image Represent. 23 (4), 622–633.

E.L. Jurdi, R., Petitjean, C., Honeine, P., Cheplygina, V., Abdallah, F., 2021. A surprisingly effective perimeter-based loss for medical image segmentation. In: Heinrich, M., Dou, Q., de Bruijne, M., Lellmann, J., Schläfer, A., Ernst, F. (Eds.), Proceedings of the Fourth Conference on Medical Imaging with Deep Learning. In: Proceedings of Machine Learning Research, vol. 143, PMLR, pp. 158–167, URL https://proceedings.mlr.press/v143/el-jurdi21a.html.

European Carotid Surgery Trialists' Collaborative Group, 1998. Randomised trial of endarterectomy for recently symptomatic carotid stenosis: final results of the MRC European carotid surgery trial (ECST). Lancet (ISSN: 0140-6736) 351 (9113), 1379–1387. http://dx.doi.org/10.1016/S0140-6736(97)09292-1, URL https://www.sciencedirect.com/science/article/pii/S0140673697092921.

Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H., 2021. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature Methods (ISSN: 1548-7105) 18 (2), 203–211. http://dx.doi.org/10.1038/s41592-020-01008-z.

Karimi, D., Salcudean, S.E., 2019. Reducing the hausdorff distance in medical image segmentation with convolutional neural networks. IEEE Trans. Med. Imaging 39 (2), 499–513.

Kervadec, H., Bahig, H., Letourneau-Guillon, L., Dolz, J., Ayed, I.B., 2021. Beyond pixel-wise supervision for segmentation: A few global shape descriptors might be surprisingly good!. In: Heinrich, M., Dou, Q., de Bruijne, M., Lellmann, J., Schläfer, A., Ernst, F. (Eds.), Proceedings of the Fourth Conference on Medical Imaging with Deep Learning. In: Proceedings of Machine Learning Research, vol. 143, PMLR, pp. 354–368, URL https://proceedings.mlr.press/v143/kervadec21a.html.

Kervadec, H., Bouchtiba, J., Desrosiers, C., Granger, E., Dolz, J., Ayed, I.B., 2019a. Boundary loss for highly unbalanced segmentation. In: International Conference on Medical Imaging with Deep Learning. PMLR, pp. 285–296.

Kervadec, H., Dolz, J., Tang, M., Granger, E., Boykov, Y., Ayed, I.B., 2019b. Constrained-CNN losses for weakly supervised segmentation. Med. Image Anal. 54, 88–99.

Kervadec, H., Dolz, J., Wang, S., Granger, E., Ayed, I.B., 2020. Bounding boxes for weakly supervised segmentation: Global constraints get close to full supervision. In: Medical Imaging with Deep Learning. PMLR, pp. 365–381.

Lin, D., Dai, J., Jia, J., He, K., Sun, J., 2016. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3159–3167.

Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A., van Ginneken, B., Sánchez, C.I., 2017. A survey on deep learning in medical image analysis. Med. Image Anal. (ISSN: 1361-8415) 42, 60–88. http://dx.doi.org/10.1016/j.media.2017.07.005, URL https://www.sciencedirect.com/science/article/pii/S1361841517301135.

Liu, F., Xu, D., Ferguson, M.S., Chu, B., Saam, T., Takaya, N., Hatsukami, T.S., Yuan, C., Kerwin, W.S., 2006. Automated in vivo segmentation of carotid plaque MRI with Morphology-Enhanced probability maps. Magnetic Resonance Med. 55 (3), 659–668.

Luo, L., Liu, S., Tong, X., Jiang, P., Yuan, C., Zhao, X., Shang, F., 2019. Carotid artery segmentation using level set method with double adaptive threshold (DATLS) on TOF-MRA images. Magnetic Resonance Imaging (ISSN: 0730-725X) 63, 123–130. http://dx.doi.org/10.1016/j.mri.2019.08.002, URL https://www.sciencedirect.com/science/article/pii/S0730725X19301791.

Luo, X., Wang, G., Song, T., Zhang, J., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T., Zhang, S., 2021. MIDeepSeg: Minimally interactive segmentation of unseen objects from medical images using deep learning. Med. Image Anal. (ISSN: 1361-8415) 72, 102102. http://dx.doi.org/10.1016/j.media.2021.102102, URL https://www.sciencedirect.com/science/article/pii/S1361841521001481.

Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M.-A., Arbel, T., Avants, B.B., Ayache, N., Buendia, P., Collins, D.L., Cordier, N., Corso, J.J., Criminisi, A., Das, T., Delingette, H., Demiralp, C., Durst, C.R., Dojat, M., Doyle, S., Festa, J., Forbes, F., Geremia, E., Glocker, B., Golland, P., Guo, X., Hamamci, A., Iftekharuddin, K.M., Jena, R., John, N.M., Konukoglu, E., Lashkari, D., Mariz, J.A., Meier, R., Pereira, S., Precup, D., Price, S.J., Raviv, T.R., Reza, S.M.S., Ryan, M., Sarikaya, D., Schwartz, L., Shin, H.-C., Shotton, J., Silva, C.A., Sousa, N., Subbanna, N.K., Szekely, G., Taylor, T.J., Thomas, O.M., Tustison, N.J., Unal, G., Vasseur, F., Wintermark, M., Ye, D.H., Zhao, L., Zhao, B., Zikic, D., Prastawa, M., Reyes, M., Van Leemput, K., 2015. The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). IEEE Trans. Med. Imaging 34 (10), 1993–2024. http://dx.doi.org/10.1109/TMI.2014.2377694.

Papandreou, G., Chen, L.-C., Murphy, K.P., Yuille, A.L., 2015. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1742–1750.

Pathak, D., Krahenbuhl, P., Darrell, T., 2015. Constrained convolutional neural networks for weakly supervised segmentation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1796–1804.

Qu, H., Wu, P., Huang, Q., Yi, J., Riedlinger, G.M., De, S., Metaxas, D.N., 2019. Weakly supervised deep nuclei segmentation using points annotation in histopathology images. In: International Conference on Medical Imaging with Deep Learning. PMLR, pp. 390–400.

Rajchl, M., Lee, M.C., Oktay, O., Kamnitsas, K., Passerat-Palmbach, J., Bai, W., Damodaram, M., Rutherford, M.A., Hajnal, J.V., Kainz, B., et al., 2016. Deepcut: Object segmentation from bounding box annotations using convolutional neural networks. IEEE Trans. Med. Imaging 36 (2), 674–683.

Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. Springer International Publishing, Cham, ISBN: 978-3-319-24574-4, pp. 234–241.

Sahasrabudhe, M., Christodoulidis, S., Salgado, R., Michiels, S., Loi, S., André, F., Paragios, N., Vakalopoulou, M., 2020. Self-supervised nuclei segmentation in histopathological images using attention. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 393–402.

Schwartz, L.H., Litière, S., de Vries, E., Ford, R., Gwyther, S., Mandrekar, S., Shankar, L., Bogaerts, J., Chen, A., Dancey, J., Hayes, W., Hodi, F.S., Hoekstra, O.S., Huang, E.P., Lin, N., Liu, Y., Therasse, P., Wolchok, J.D., Seymour, L., 2016. RECIST 1.1-update and clarification: From the RECIST committee. 2016/05/14 Euro. J. Cancer (Oxford, England : 1990) (ISSN: 1879-0852) 62, 132–137. http://dx.doi.org/10.1016/j.ejca.2016.03.081, URL https://pubmed.ncbi.nlm.nih.gov/27189322, 27189322[pmid].

Tang, M., Perazzi, F., Djelouah, A., Ben Ayed, I., Schroers, C., Boykov, Y., 2018. On regularized losses for weakly-supervised cnn segmentation. In: Proceedings of the European Conference on Computer Vision. ECCV, pp. 507–522.

Truijman, M., Kooi, M.E., Van Dijk, A., de Rotte, A., van der Kolk, A., Liem, M., Schreuder, F., Boersma, E., Mess, W., van Oostenbrugge, R.J., et al., 2014. Plaque At RISK (PARISK): prospective multicenter study to improve diagnosis of high-risk carotid plaques. Int. J. Stroke 9 (6), 747–754.

van Engelen, A., van Dijk, A.C., Truijman, M.T.B., van't Klooster, R., van Opbroek, A., van der Lugt, A., Niessen, W.J., Kooi, M.E., de Bruijne, M., 2015. Multi-center MRI carotid plaque component segmentation using feature normalization and transfer learning. IEEE Trans. Med. Imaging 34 (6), 1294–1305. http://dx.doi.org/10.1109/TMI.2014.2384733.

Wang, G., Li, W., Zuluaga, M.A., Pratt, R., Patel, P.A., Aertsen, M., Doel, T., David, A.L., Deprest, J., Ourselin, S., Vercauteren, T., 2018. Interactive medical image segmentation using deep learning with image-specific fine tuning. IEEE Trans. Med. Imaging 37 (7), 1562–1573. http://dx.doi.org/10.1109/TMI.2018.2791721.

Wu, J., Xin, J., Yang, X., Sun, J., Xu, D., Zheng, N., Yuan, C., 2019. Deep morphology aided diagnosis network for segmentation of carotid artery vessel wall and diagnosis of carotid atherosclerosis on black-blood vessel wall MRI. Med. Phys. 46 (12), 5544–5561.

Yang, H., Deng, R., Lu, Y., Zhu, Z., Chen, Y., Roland, J.T., Lu, L., Landman, B.A., Fogo, A.B., Huo, Y., 2020. CircleNet: Anchor-free glomerulus detection with circle representation. In: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racoceanu, D., Joskowicz, L. (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2020. Springer International Publishing, Cham, ISBN: 978-3-030-59719-1, pp. 35–44.

Zhang, Q., Qiao, H., Dou, J., Sui, B., Zhao, X., Chen, Z., Wang, Y., Chen, S., Lin, M., Chiu, B., Yuan, C., Li, R., Chen, H., 2019. Plaque components segmentation in carotid artery on simultaneous non-contrast angiography and intraplaque hemorrhage imaging using machine learning. Magnetic Resonance Imaging (ISSN: 0730-725X) 60, 93–100. http://dx.doi.org/10.1016/j.mri.2019.04.001, URL https://www.sciencedirect.com/science/article/pii/S0730725X18306313.

Zhao, X., Li, R., Hippe, D.S., Hatsukami, T.S., Yuan, C., Investigators, C.-I., 2017. Chinese Atherosclerosis Risk Evaluation (CARE II) study: a novel cross-sectional, multicentre study of the prevalence of high-risk atherosclerotic carotid plaque in Chinese patients with ischaemic cerebrovascular events—design and rationale. Stroke Vasc. Neurol. (ISSN: 2059-8688) 2 (1), 15–20. http://dx.doi.org/10.1136/svn-2016-000053, arXiv:https://svn.bmj.com/content/2/1/15.full.pdf.

Zhu, C., Wang, X., Teng, Z., Chen, S., Huang, X., Xia, M., Mao, L., Bai, C., 2021. Cascaded residual U-net for fully automatic segmentation of 3D carotid artery in high-resolution multi-contrast MR images. Phys. Med. Biol. 66 (4), 045033.