

Differentiable Boundary Point Extraction for Weakly Supervised Star-shaped Object Segmentation

Robin Camarasa^{1,2*}

R.CAMARASA@ERASMUSMC.NL

Hoel Kervadec^{1,2}

H.KERVADEC@ERASMUSMC.NL

Daniel Bos^{2,3}

D.BOS@ERASMUSMC.NL

Marleen de Bruijne^{1,2,4}

MARLEEN.DEBRUIJNE@ERASMUSMC.NL

¹ *Biomedical Imaging Group Rotterdam, Erasmus MC, Rotterdam, The Netherlands*

² *Department of Radiology and Nuclear Medicine, Erasmus MC, Rotterdam, The Netherlands*

³ *Department of Epidemiology, Erasmus MC, Rotterdam, The Netherlands*

⁴ *Department of Computer Science, University of Copenhagen, Denmark*

Editors: Under Review for MIDL 2022

Abstract

Although Deep Learning is the new gold standard in medical image segmentation, the annotation burden limits its expansion to clinical practice. We also observe a mismatch between annotations required by deep learning methods designed with pixel-wise optimization in mind and clinically relevant annotations designed for biomarkers extraction (diameters, counts, etc.). Our study proposes a first step toward bridging this gap, optimizing vessel segmentation based on its diameter annotations. To do so we propose to extract boundary points from a star-shaped segmentation in a differentiable manner. This differentiable extraction allows reducing annotation burden as instead of the pixel-wise segmentation only the two annotated points required for diameter measurement are used for training the model. Our experiments show that training based on diameter is efficient; produces state-of-the-art weakly supervised segmentation; and performs reasonably compared to full supervision.

Our code is publicly available:

<https://gitlab.com/radiology/aim/carotid-artery-image-analysis/diameter-learning>

Keywords: Image segmentation ; weak annotations ; Carotid artery stenosis

1. Introduction

In the recent years, huge steps have been made in the direction of automated medical image segmentation as a consequence of, on the one hand, the development of new and powerful deep learning models, and, on the other hand, the increasing availability of open access to large, high-quality (e.g. fully segmented and accurate) datasets for model training. However, as each task is unique, data-scarcity remains an issue for many applications. In most applications full segmentations are not made during clinical practice and therefore are rarely available. However, in many cases, simpler annotations are made by the clinicians, not with a segmentation goal in mind, but rather to estimate the relevant imaging bio-markers.

* Corresponding author

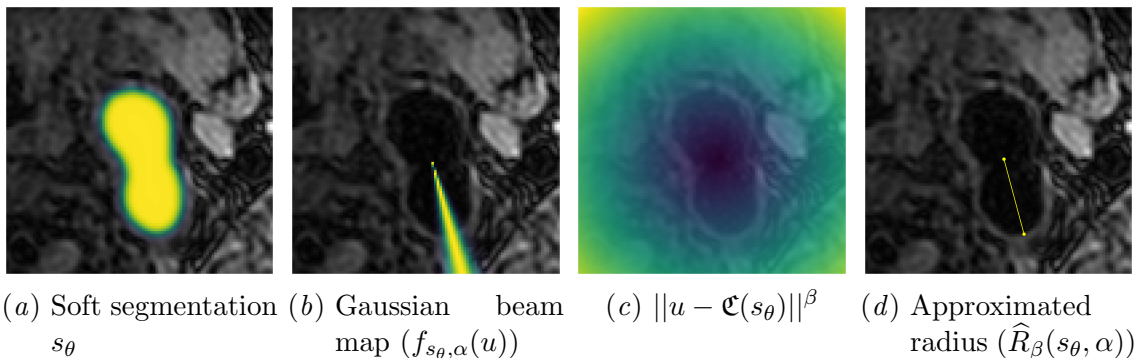


Figure 1: Overview of the different steps of our method—overlayed on an MRI—going from continuous probabilities (a) to a boundary point (d). Best viewed in colors.

For instance, in the case of carotid artery atherosclerosis¹, clinicians measure the diameter of the lumen at different places in the artery, enabling them to quickly assess the degree of stenosis (i.e. the local *narrowing* of the artery, due to the wall *thickening*). Although, from a clinical perspective it is widely acknowledged that a full segmentation of the lumen would be more informative, and could enable more thorough analysis, it is too time consuming to put in practice.

Yet, the wealth of existing diameter measurements, could be used to train a segmentation network which could ultimately provide full-lumen segmentation and allow to train models on larger and more varied datasets without extra annotation costs. This type of task-specific weak label has not, to the best of our knowledge, been used in the past. As such, existing and popular weakly supervised methods cannot be used directly and a new and efficient way to use those annotations is required.

Our contributions can be summarized as follows:

- we propose a differentiable module, that for a given angle measure the radius of a predicted segmentation;
- we then comprehensively evaluate the proposed method with 3 runs of 4-fold cross-validation and compare to the relevant existing methods from the literature.

2. Related work

Weakly supervised segmentation has received an increased attention in the recent years, and 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) and even classification labels (image tags) (Pathak et al., 2015; Kervadec et al., 2019; Sahasrabudhe

1. The thickening of the vessel wall of the carotid artery, an important risk factor for ischemic stroke (World Health Organization, 2011) and a leading cause of death and disability worldwide (World Health Organization, 2014).

et al., 2020). Most of those methods are “generic”, in the sense that they can be applied to a variety of tasks without further modification.

However, when strong priors on the segmented shape exists, some weakly supervised methods can be tailored to efficiently exploit this information and compensate the lack of full-annotations. For instance, Sahasrabudhe et al. (2020) and Qu et al. (2019) exploit the structure of deep nuclei segmentation by either learning 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 “draw” a path between the extreme points annotations, by minimizing a geodesic distance. In a multi-instance object detection problem, 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 as label. 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 diameter information but ignoring the exact boundary of the object (as it simply predicts a circle). As such, to the best of our knowledge, there exists no method in the current literature that could use the maximum diameter information while segmenting the object precisely.

3. Method

Let us denote² an input image $I : \Omega \rightarrow \mathbb{R}$, with $\Omega \in \mathbb{R}^2$ a 2D image space, for which we want to predict its corresponding segmentation $y : \Omega \rightarrow \{0, 1\}$. Due to the supervision setting, we do not have access to y at training time, and must instead rely on the diameter annotations: $(l, l') \in \Omega^2$, which are simply two points placed on the object boundary. Those points enable us to deduce more information about the object: i) the location of its centroid, ii) its maximum diameter.

To achieve our goal of segmenting the carotid artery, we proceed in different steps: 3.1) locating the centroid of the segmentation, 3.2) extracting points on the boundary of the predicted segmentation in a differentiable way, 3.3) deriving the diameter of the predicted segmentation, and 3.4) creating a combined loss using the aforementioned modules. The first three steps are illustrated in Figure 2.

3.1. Location of the vessel

As multiple blood vessels can be present in a single image, the first step is to locate the segmented artery; to “anchor” the segmentation in the correct place during training. As Kervadec et al. (2021) did, the centroid can be efficiently computed as a probability-weighted average of pixel coordinates³:

$$\mathfrak{C}(s_\theta) = \frac{1}{\int_{u \in \Omega} s_\theta(u) du} \int_{u \in \Omega} u \cdot s_\theta(u) du, \quad (1)$$

2. For convenience, all mathematical notation is summarized in Appendix A, though each symbol is introduced in time in the manuscript.

3. In this paper, we use an integral notation, for consistency with the rest of the formulation. The implementation will eventually discretize it as a sum.

where $s_\theta : \Omega \rightarrow [0, 1]$ are the predicted probabilities. $\mathfrak{C}(s_\theta)$ is illustrated in Figure 2(a).

3.2. Boundary points extraction

With the centroid information, we can derive the boundary coordinates of the segmentation, for different angles $\alpha \in [0, 2\pi[$. This is done in several sub-steps.

First, a Gaussian beam map (shown in Figure 1(b)) restricts the computation to a continuous area around the angle α :

$$f_{s_\theta, \alpha}(u) = g_\alpha(\Phi_{s_\theta}(u)_\gamma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\Phi_{s_\theta}(u)_\gamma - \alpha)^2}{2\sigma^2}\right), \quad (2)$$

where $\Phi_{s_\theta} : \Omega \rightarrow \mathbb{R}^+ \times [0, 2\pi[$ is a function to go from Cartesian to polar coordinates, σ a hyper-parameter, controlling the “width” of the beam. $\Phi_{s_\theta}(u)_\gamma$ is the angular coordinate. The Gaussian beam-map can then be embedded into a weighted *average distance to the centroid* of the soft segmentation:

$$\mathcal{I}_\beta(s_\theta, \alpha) = \int_{u \in \Omega} \underbrace{f_{s_\theta, \alpha}(u)}_{\text{beam}} s_\theta(u) \underbrace{\|u - \mathfrak{C}(s_\theta)\|_2^\beta}_{\text{dist. to centroid}} du, \quad (3)$$

where β is an exponent balancing the contribution of the *distance to the centroid*. It can be shown—and this is the main theoretical result of this paper, see Appendix B—that this weighted average is a *differentiable* approximation of the exact radius:

$$\mathcal{I}_\beta(s_\theta, \alpha) \approx \frac{\mathfrak{R}(s_\theta, \alpha)^{\beta+2}}{\beta + 2}. \quad (4)$$

This enables the definition of a first approximate radius function $\widehat{\mathfrak{R}}_\beta$:

$$\widehat{\mathfrak{R}}_\beta(s_\theta, \alpha) = ((\beta + 2) \mathcal{I}_\beta(s_\theta, \alpha))^{\frac{1}{\beta+2}}. \quad (5)$$

Notice that this approximate still depends on a parameter β . As multiple soft segmentations would result in the same radius estimate for a single β , we average different radius estimates over several β to obtain a stronger final radius estimation as:

$$\widehat{\mathfrak{R}}(s_\theta, \alpha) = \mathbb{E}_{\beta \in B} \left(\widehat{\mathfrak{R}}_\beta(s_\theta, \alpha) \right), \quad (6)$$

with $B \subset \mathbb{R}^+$ a discrete set of β . Finally, we can go back to the Cartesian space with the function $\Phi_{s_\theta}^{-1}(\cdot, \alpha) : \mathbb{R}^+ \times [0, 2\pi[\rightarrow \Omega$ that reverses the initial coordinates transform Φ_{s_θ} for the angle α used in Eq. (2). Notice that this radius estimation is fully differentiable, and as such, can be back-propagated over.

3.3. Diameter extraction

The diameter estimate is then defined as the maximum distance between the different boundary coordinates, for a discrete set Γ of angles:

$$\widehat{\mathfrak{D}}(s_\theta) = \max \left\{ \left\| \Phi_{s_\theta}^{-1} \left(\widehat{\mathfrak{R}}(s_\theta, \alpha), \alpha \right) - \Phi_{s_\theta}^{-1} \left(\widehat{\mathfrak{R}}(s_\theta, \alpha'), \alpha' \right) \right\|_2, \right. \\ \left. \forall (\alpha, \alpha') \in \Gamma^2 \right\}. \quad (7)$$

3.4. Losses, regularizer, and final model

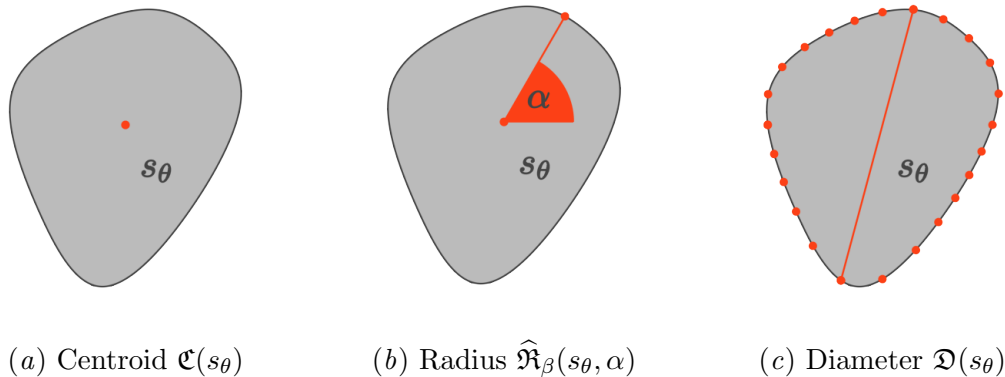


Figure 2: Schematic representation of the differentiable module

The centroid and diameter can be learned with a mean squared error, using ground truth centroid ($\frac{l+l'}{2}$) and diameter ($\|l-l'\|_2$) as reference:

$$\mathcal{L}_{\mathfrak{C}}(\theta) = \left\| \mathfrak{C}(s_\theta) - \frac{l+l'}{2} \right\|_2^2, \quad (8)$$

$$\mathcal{L}_{\mathfrak{D}}(\theta) = \left\| \widehat{\mathfrak{D}}(s_\theta) - \|l-l'\|_2 \right\|_2^2. \quad (9)$$

This provides most of the supervision, but, as we will show in the experiments section, additional regularization is desirable. Imposing a star-shape prior (favoring continuous segmentation with a sharp boundary) is an effective way to do it. This translate into a minimization of the variance of the radius estimate for different β :

$$\mathcal{R}(\theta) = \mathbb{E}_{\alpha \in \Gamma} \left(\text{Var}_{\beta \in B} \left(\widehat{\mathfrak{R}}_\beta(s_\theta, \alpha) \right) \right). \quad (10)$$

The final combined loss is therefore:

$$\mathcal{L}(\theta) = \lambda \mathcal{L}_{\mathfrak{C}}(\theta) + \mu \mathcal{L}_{\mathfrak{D}}(\theta) + \nu \mathcal{R}(\theta), \quad (11)$$

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

4. Experiments

4.1. Data

We use the data from the CARE II study (Zhao et al., 2017), where the 24 enrolled patients had a recent ischaemic stroke or transient ischaemic attack. 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 average in 12.1% of the slices of a scan. Based on those segmented slices we simulate diameter annotations by picking for each (l, l') such as the points are the furthest away

from each others. This result in 2151 annotated 2D slices over the whole dataset. At training and evaluation time all 2D slices are cut in half to have only one carotid artery per image.

4.2. Baselines and ablation study

We compare our method to InExtremIS (Dorent et al., 2021) and CircleNet (Yang et al., 2020), as they are the most relevant and related weakly supervised methods to our problem. To perform a fair comparison across all methods, the additional CRF loss (Tang et al., 2018) used by InExtremeIS is not included; as all methods could benefit from it. We believe that this enables a more exact analysis of each methods performances and limitations. For CircleNet, the input and output spaces have been tweaked to have a similar dimension. This enables us to use the same network architecture for all methods. 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. (11).

4.3. Implementation details

All methods are build on top of the base U-Net architecture (Ronneberger et al., 2015)⁴ and are trained with the same ADAM optimizer (e.g. learning rate 10^{-4} , $\beta_1 = 0.9$, $\beta_2 = 0.99$) for 600 epochs. Our loss components (Equations (8), (9) and (10)) do not require any modification of the network architecture or training regiment, and are implemented as direct losses. We use $|\Gamma| = 24$ radii equally spread between $[0, 2\pi[$ and $B = \{0, 1\}$ as β exponents. λ , μ , ν and σ have been empirically set to 10, 100, 1 and 0.15 experimenting on a private dataset. For all methods the final segmentation is the largest connected component of the network output thresholded at a value 0.5 as having multiple components is not realistic and would result in a systematic over estimation of the diameter. Our publicly available code⁵ contains the implementations of both our methods and the different baselines.

4.4. Metrics and evaluation

All methods were trained and evaluated over a 4-folds cross-validation (2 folds for training, 1 fold for validation and 1 fold for testing), repeated three times with different random seeds. Evaluation is performed with dice score (DSC), Hausdorff distance⁶ (HD) and absolute diameter error (ADE); we report the per-patient average over the testing set (over each run). We also perform a two-sided Wilcoxon signed-rank test, with a level of significance of 0.05 to determine if a run of the cross-validation is significantly different from our method.

5. Results

We display the metrics distribution in Figure 3, for each run of the cross-validation. In Table 1, we report the mean (and standard deviation) of the medians of those distributions, computed over the different runs of the cross-validation.

4. With the official Monai implementation as starting point: <https://docs.monai.io/en/0.7>.

5. <https://gitlab.com/radiology/aim/carotid-artery-image-analysis/diameter-learning>

6. In the case of an empty foreground (no prediction), we set the Hausdorff distance to the diagonal of the image, i.e. 272.4 mm.

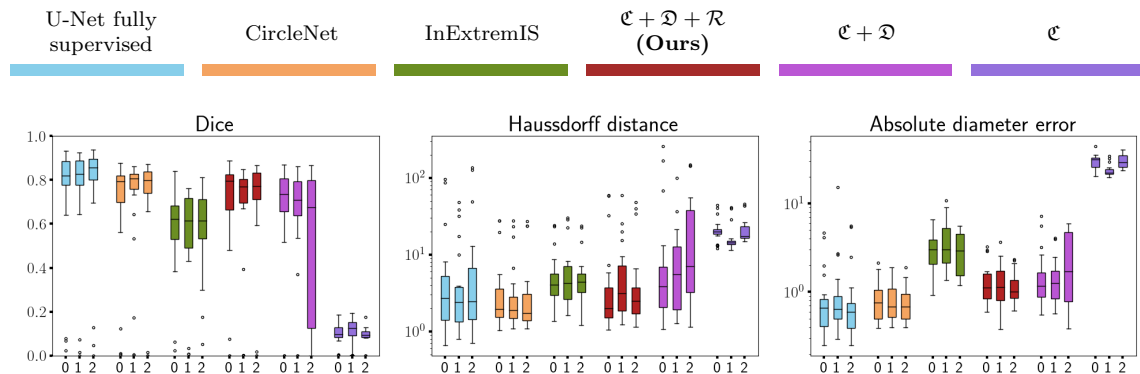


Figure 3: Boxplots of ours metrics, averaged per patient. Each box represents a run of the cross-validation, and go from Q1 to Q3 while displaying the median. The whiskers extends from the box by 1.5 time the inter-quartile range.

Table 1: Mean and standard deviation over the different runs of the medians of the metrics averaged per patient. The superscript highlight runs that are statistically different from our proposed method.

Method	DSC (%) \uparrow	HD (mm) \downarrow	ADE (mm) \downarrow
U-Net fully supervised	0.835(0.016) ^{0,1,2}	2.53(0.13)	0.64(0.03) ^{0,1,2}
CircleNet (Yang et al., 2020)	0.799(0.005) ¹	1.85(0.08) ^{1,2}	0.71(0.04) ^{0,1,2}
InExtremIS (Dorent et al., 2021)	0.617(0.003) ^{0,1,2}	4.27(0.16)	2.98(0.02) ^{0,1,2}
$\mathfrak{C} + \mathfrak{D} + \mathfrak{R}$ (Ours)	0.780(0.012)	2.54(0.47)	1.09(0.05)
$\mathfrak{C} + \mathfrak{D}$	0.706(0.025) ^{0,1,2}	5.47(1.29) ^{0,1,2}	1.38(0.24) ^{1,2}
\mathfrak{C}	0.105(0.015) ^{0,1,2}	17.14(2.24) ^{0,1,2}	27.64(3.95) ^{0,1,2}

We can quickly notice that the best weakly supervised methods (CircleNet and ours) are fairly close to full-supervision performances (3.6% and 5.5% differences, respectively). However, notice that the difference between CircleNet and our method, despite the slightly higher performances for CircleNet (partly explained by the approximative circularity of most segmented arteries, see Appendix D, Figure 6), is not statistically different on each run. The ablation study demonstrates the usefulness of each component of the loss, as removing the regularizer \mathfrak{R} induces a 8% drop in performances, while supervising only the centroid \mathfrak{C} gives very poor performances.

Moreover, the qualitative results from Figure 4 highlight a key difference between CircleNet and our method: as CircleNet predict only perfect circles, the segmentation precision decreases when shape-complexity increases (Appendix D, Figure 7). An example is in the second row of Figure 4, where the shape of the carotid in its bifurcation is lost.

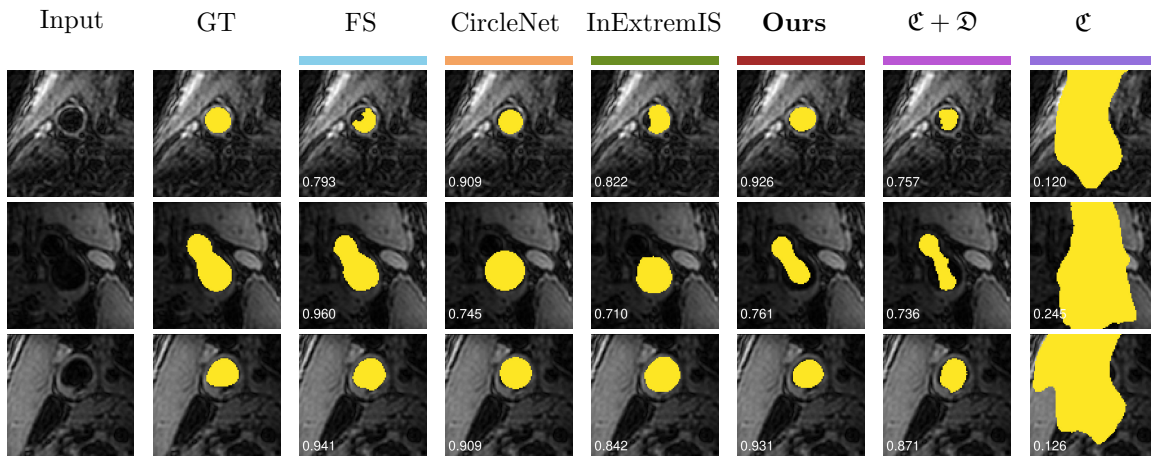


Figure 4: Obtained carotid artery segmentation for the different methods and different examples. The bottom left value indicates the Dice score.

6. Discussion

The experiments have shown the strong performances (both quantitatively and qualitatively) of our method, over existing weakly supervised methods. Although CircleNet performed better than our proposed method it does not propose a viable alternative as it does not fully capture the complexity of the shape of the carotid artery bifurcation. However, a more comprehensive evaluation remains to be done. Most notably, we had to simulate the diameter annotations: it would be interesting in a future work to compare clinician annotations to simulated annotations. Indeed, the clinician annotations might not always be *exactly* at the widest diameter of the image, but this slight imprecision might prove (or not) negligible. Another limitation of the current study came from the data heterogeneity: all methods failed in the (under-represented) subset of images ($n = 3$) that had a different spatial resolution.

While the differentiable boundary extraction module in this paper is used for diameter optimization, we are not exploiting it to its full potential: it could be used to model and supervise more complex shapes (Appendix C , Figure 5), 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 method to locate the centroid and boundary points of an object. This powerful method can then be used to train a neural network, supervising the object shape around its centroid. This provides an elegant way to re-use existing clinical annotations. We validated the method on the task of carotid artery segmentation, using only diameter annotations for training, and showed state of the art weakly supervised segmentation performance.

References

- Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *European conference on computer vision*, pages 549–565. Springer, 2016.
- Reuben Dorent, Samuel Joutard, Jonathan Shapey, Aaron Kujawa, Marc Modat, Sébastien Ourselin, and Tom Vercauteren. Inter extreme points geodesics for end-to-end weakly supervised image segmentation. In Marleen de Bruijne, Philippe C. Cattin, Stéphane Cotin, Nicolas Padoy, Stefanie Speidel, Yefeng Zheng, and Caroline Essert, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, pages 615–624, Cham, 2021. Springer International Publishing. ISBN 978-3-030-87196-3.
- Hoel Kervadec, Jose Dolz, Meng Tang, Eric Granger, Yuri Boykov, and Ismail Ben Ayed. Constrained-cnn losses for weakly supervised segmentation. *Medical image analysis*, 54: 88–99, 2019.
- Hoel Kervadec, Jose Dolz, Shanshan Wang, Eric Granger, and Ismail Ben Ayed. Bounding boxes for weakly supervised segmentation: Global constraints get close to full supervision. In *Medical Imaging with Deep Learning*, pages 365–381. PMLR, 2020.
- Hoel Kervadec, Houda Bahig, Laurent Letourneau-Guillon, Jose Dolz, and Ismail Ben Ayed. Beyond pixel-wise supervision for segmentation: A few global shape descriptors might be surprisingly good! In Mattias Heinrich, Qi Dou, Marleen de Bruijne, Jan Lellmann, Alexander Schläfer, and Floris Ernst, editors, *Proceedings of the Fourth Conference on Medical Imaging with Deep Learning*, volume 143 of *Proceedings of Machine Learning Research*, pages 354–368. PMLR, 07–09 Jul 2021. URL <https://proceedings.mlr.press/v143/kervadec21a.html>.
- Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3159–3167, 2016.
- Xiangde Luo, Guotai Wang, Tao Song, Jingyang Zhang, Michael Aertsen, Jan Deprest, Sebastien Ourselin, Tom Vercauteren, and Shaoting Zhang. Mideepseg: Minimally interactive segmentation of unseen objects from medical images using deep learning. *Medical Image Analysis*, 72:102102, 2021. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2021.102102>. URL <https://www.sciencedirect.com/science/article/pii/S1361841521001481>.
- Bjoern H. Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, Levente Lencsés, Elizabeth Gerstner, Marc-André Weber, Tal Arbel, Brian B. Avants, Nicholas Ayache, Patricia Buendia, D. Louis Collins, Nicolas Cordier, Jason J. Corso, Antonio Criminisi, Tilak Das, Hervé Delingette, Çağatay Demiralp, Christopher R. Durst, Michel Dojat, Senan Doyle, Joana Festa, Florence Forbes, Ezequiel Geremia, Ben Glocker, Polina Golland, Xiaotao Guo, Andac Hamamci, Khan M. Iftekharruddin, Raj Jena, Nigel M. John, Ender Konukoglu, Danial Lashkari, José António Mariz, Raphael

- Meier, Sérgio Pereira, Doina Precup, Stephen J. Price, Tammy Riklin Raviv, Syed M. S. Reza, Michael Ryan, Duygu Sarikaya, Lawrence Schwartz, Hoo-Chang Shin, Jamie Shotton, Carlos A. Silva, Nuno Sousa, Nagesh K. Subbanna, Gabor Szekely, Thomas J. Taylor, Owen M. Thomas, Nicholas J. Tustison, Gozde Unal, Flor Vasseur, Max Wintermark, Dong Hye Ye, Liang Zhao, Binsheng Zhao, Darko Zikic, Marcel Prastawa, Mauricio Reyes, and Koen Van Leemput. The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). *IEEE Transactions on Medical Imaging*, 34(10):1993–2024, 2015. doi: 10.1109/TMI.2014.2377694.
- George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015.
- Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1796–1804, 2015.
- Hui Qu, Pengxiang Wu, Qiaoying Huang, Jingru Yi, Gregory M Riedlinger, Subhajyoti De, and Dimitris N Metaxas. Weakly supervised deep nuclei segmentation using points annotation in histopathology images. In *International Conference on Medical Imaging with Deep Learning*, pages 390–400. PMLR, 2019.
- Martin Rajchl, Matthew CH Lee, Ozan Oktay, Konstantinos Kamnitsas, Jonathan Passerat-Palmbach, Wenjia Bai, Mellisa Damodaram, Mary A Rutherford, Joseph V Hajnal, Bernhard Kainz, et al. Deepcut: Object segmentation from bounding box annotations using convolutional neural networks. *IEEE transactions on medical imaging*, 36(2):674–683, 2016.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24574-4.
- Mihir Sahasrabudhe, Stergios Christodoulidis, Roberto Salgado, Stefan Michiels, Sherene Loi, Fabrice André, Nikos Paragios, and Maria Vakalopoulou. Self-supervised nuclei segmentation in histopathological images using attention. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 393–402. Springer, 2020.
- Lawrence H. Schwartz, Saskia Litière, Elisabeth de Vries, Robert Ford, Stephen Gwyther, Sumithra Mandrekar, Lalitha Shankar, Jan Bogaerts, Alice Chen, Janet Dancey, Wendy Hayes, F. Stephen Hodi, Otto S. Hoekstra, Erich P. Huang, Nancy Lin, Yan Liu, Patrick Therasse, Jedd D. Wolchok, and Lesley Seymour. Recist 1.1-update and clarification: From the recist committee. *European journal of cancer (Oxford, England : 1990)*, 62: 132–137, Jul 2016. ISSN 1879-0852. doi: 10.1016/j.ejca.2016.03.081. URL <https://pubmed.ncbi.nlm.nih.gov/27189322>. 27189322[pmid].

- Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On regularized losses for weakly-supervised cnn segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 507–522, 2018.
- Guotai Wang, Wenqi Li, Maria A. Zuluaga, Rosalind Pratt, Premal A. Patel, Michael Aertsen, Tom Doel, Anna L. David, Jan Deprest, Sébastien Ourselin, and Tom Vercauteren. Interactive Medical Image Segmentation Using Deep Learning With Image-Specific Fine Tuning. *IEEE Transactions on Medical Imaging*, 37(7):1562–1573, 2018. doi: 10.1109/TMI.2018.2791721.
- World Health Organization. Global atlas on cardiovascular disease prevention and control: published by the World Health Organization in collaboration with the World Heart Federation and the World Stroke Organization. 2011.
- World Health Organization. *Global status report on noncommunicable diseases 2014*. Number WHO/NMH/NVI/15.1. World Health Organization, 2014.
- Haichun Yang, Ruining Deng, Yuzhe Lu, Zheyu Zhu, Ye Chen, Joseph T. Roland, Le Lu, Bennett A. Landman, Agnes B. Fogo, and Yuankai Huo. Circlenet: Anchor-free glomerulus detection with circle representation. In Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, Diana Mateus, Maria A. Zuluaga, S. Kevin Zhou, Daniel Racoceanu, and Leo Joskowicz, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, pages 35–44, Cham, 2020. Springer International Publishing. ISBN 978-3-030-59719-1.
- Xihai Zhao, Rui Li, Daniel S Hippe, Thomas S Hatsukami, Chun Yuan, and CARE-II Investigators. 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 and Vascular Neurology*, 2(1):15–20, 2017. ISSN 2059-8688. doi: 10.1136/svn-2016-000053. URL <https://svn.bmj.com/content/2/1/15>.

Appendix A. List of Symbols

$u \in \Omega \subset \mathbb{R}^2$	coordinates in the Cartesian image space;
$I : \Omega \rightarrow \mathbb{R}$	input image;
$s_\theta : \Omega \rightarrow [0, 1]$	softmax predictions for input I and parameters θ ;
$(l, l') \in \Omega^2$	simulated diameter annotations;
$\mathfrak{C} : (\Omega \rightarrow [0, 1]) \rightarrow \Omega$	function computing the centroid of a (predicted) segmentation;
$(r, \gamma) \in \Lambda_\theta = \mathbb{R}^+ \times [0, 2\pi[$	polar coordinates, centered around $\mathfrak{C}(s_\theta)$;
$\Phi_{s_\theta} : \Omega \rightarrow \Lambda_\theta$	function going from Cartesian to polar coordinates;
$\alpha \in \Gamma \subset [0, 2\pi[$	some angle out of a discrete set;
$\beta \in B \subset \mathbb{R}^+$	some exponent out of a discrete set;
$\mathfrak{R} : (\Omega \rightarrow \{0, 1\}) \times [0, 2\pi[\rightarrow \mathbb{R}^+$	exact radius of a discrete segmentation;
$\widehat{\mathfrak{R}}_\beta : (\Omega \rightarrow [0, 1]) \times [0, 2\pi[\rightarrow \mathbb{R}^+$	approximate radius of continuous probabilities with param β ;
$\widehat{\mathfrak{R}} : (\Omega \rightarrow [0, 1]) \times [0, 2\pi[\rightarrow \mathbb{R}^+$	final approximate radius of continuous probabilities;
$\widehat{\mathfrak{D}} : (\Omega \rightarrow [0, 1]) \rightarrow \mathbb{R}^+$	approximated diameter for continuous probabilities;
$(\lambda, \mu, \nu) \in \mathbb{R}^3$	hyper-parameters balancing the combined loss

Appendix B. Radius derivations

Let's derive the integral of Equation (3):

$$\mathcal{I}_\beta(s_\theta, \alpha) = \int_{u \in \Omega} f_{s_\theta, \alpha}(u) s_\theta(u) \|u - \mathfrak{C}(s_\theta)\|_2^\beta du. \quad (12)$$

We apply the following change of variable: $(r, \gamma) = \Phi_{s_\theta}(u)$

$$\mathcal{I}_\beta(s_\theta, \alpha) = \int \int_{(r, \gamma) \in \Lambda_\theta} f_{s_\theta, \alpha}(\Phi_{s_\theta}^{-1}(r, \gamma)) s_\theta(\Phi_{s_\theta}^{-1}(r, \gamma)) r^\beta r dr d\gamma \quad (13)$$

$$= \int \int_{(r, \gamma) \in \Lambda_\theta} g_\alpha(\gamma) s_\theta(\Phi_{s_\theta}^{-1}(r, \gamma)) r^{\beta+1} dr d\gamma \quad (14)$$

$$= \int_{\gamma \in [0, 2\pi]} \int_{r \in \mathbb{R}^+} g_\alpha(\gamma) \mathbf{1}_{x \leq \mathfrak{R}(s_\theta, \gamma)}(r) r^{\beta+1} dr d\gamma \quad (15)$$

$$= \int_{\gamma \in [0, 2\pi]} \int_{r=0}^{\mathfrak{R}(s_\theta, \gamma)} g_\alpha(\gamma) r^{\beta+1} dr d\gamma \quad (16)$$

$$= \frac{1}{\beta + 2} \int_{\gamma \in [0, 2\pi]} g_\alpha(\gamma) \mathfrak{R}(s_\theta, \gamma)^{\beta+2} d\gamma. \quad (17)$$

Under the following assumptions:

$$\begin{aligned} \int_{\gamma \in \mathbb{R}} g_\alpha(\gamma) d\gamma &\approx \int_{\gamma \in [\alpha-3\sigma, \alpha+3\sigma]} g_\alpha(\gamma) d\gamma \\ &\approx 1 \\ \mathfrak{R}(s_\theta, \gamma) &\approx \mathfrak{R}(s_\theta, \alpha) \quad \forall \gamma \in [\alpha - 3\sigma, \alpha + 3\sigma], \end{aligned}$$

$$\mathcal{I}_\beta(s_\theta, \alpha) \approx \frac{1}{\beta + 2} \int_{\gamma \in [\alpha-3\sigma, \alpha+3\sigma]} g_\alpha(\gamma) \mathfrak{R}(s_\theta, \alpha)^{\beta+2} d\gamma \quad (18)$$

$$\approx \frac{\mathfrak{R}(s_\theta, \alpha)^{\beta+2}}{\beta + 2}. \quad (19)$$

Finally we obtain the radius estimate of Equation (5)

$$\widehat{\mathfrak{R}}_\beta(s_\theta, \alpha) = ((\beta + 2)\mathcal{I}_\beta(s_\theta, \alpha))^{\frac{1}{\beta+2}}. \quad (20)$$

Appendix C. Synthetic examples




Figure 5: Application of the differentiable boundary point extraction applied to randomly generated star-shaped objects (green: generated star-shape object, red: extracted boundary points, blue: bounding circle).

Appendix D. Shape analysis

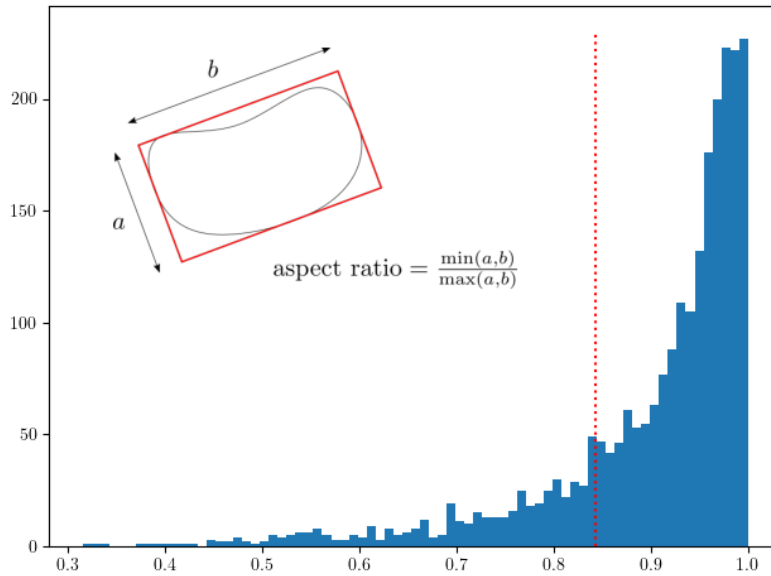


Figure 6: Distribution of the slice-wise aspect ratios of the carotid artery segmentations (the red dotted line represents the 80th percentile).

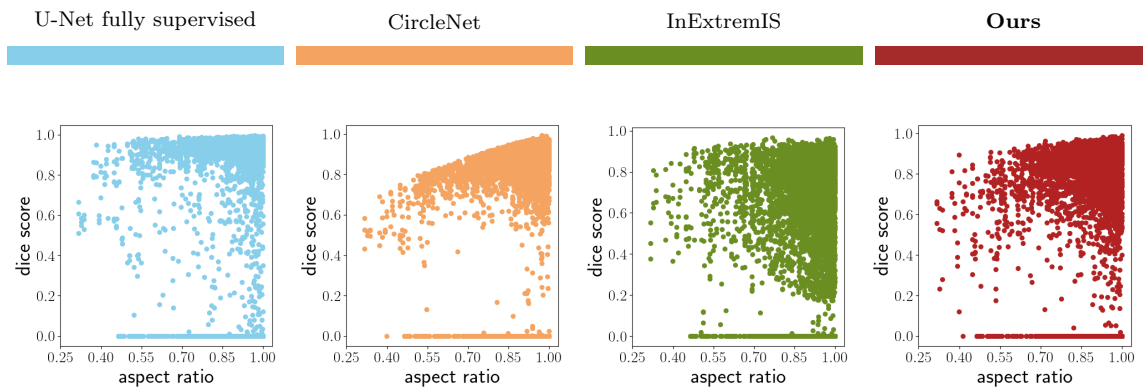


Figure 7: Slice-wise Dice as a function of the slice-wise aspect ratio per method. (Results of all runs of the cross-validation)