



**HAL**  
open science

# Guidelines for cerebrovascular segmentation: Managing imperfect annotations in the context of semi-supervised learning

Pierre Rougé, Pierre-Henri Conze, Nicolas Passat, Odysée Merveille

## ► To cite this version:

Pierre Rougé, Pierre-Henri Conze, Nicolas Passat, Odysée Merveille. Guidelines for cerebrovascular segmentation: Managing imperfect annotations in the context of semi-supervised learning. *Computerized Medical Imaging and Graphics*, 2025, 119, pp.102474. 10.1016/j.compmedimag.2024.102474 . hal-04819764

**HAL Id: hal-04819764**

**<https://hal.science/hal-04819764v1>**

Submitted on 5 Dec 2024

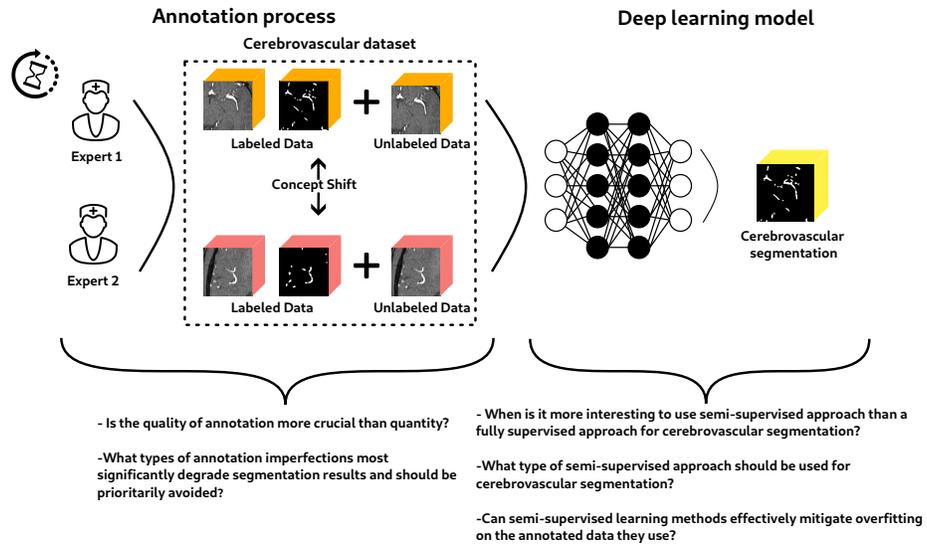
**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Graphical Abstract

## Guidelines for Cerebrovascular Segmentation: Managing Imperfect Annotations in the Context of Semi-Supervised Learning

Pierre Roug , Pierre-Henri Conze, Nicolas Passat, Odyssee Merveille



# Guidelines for Cerebrovascular Segmentation: Managing Imperfect Annotations in the Context of Semi-Supervised Learning

Pierre Rougé<sup>a,b</sup>, Pierre-Henri Conze<sup>c</sup>, Nicolas Passat<sup>a</sup>, Odyssee Merveille<sup>b</sup>

<sup>a</sup>Université de Reims Champagne Ardenne, CRESTIC, Reims, France,

<sup>b</sup>Univ Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, CREATIS, Lyon, France,

<sup>c</sup>IMT Atlantique, LaTIM UMR 1101, Inserm, Brest, France,

---

## Abstract

Segmentation in medical imaging is an essential and often preliminary task in the image processing chain, driving numerous efforts towards the design of robust segmentation algorithms. Supervised learning methods achieve excellent performances when fed with a sufficient amount of labeled data. However, such labels are typically highly time-consuming, error-prone and expensive to produce. Alternatively, semi-supervised learning approaches leverage both labeled and unlabeled data, and are very useful when only a small fraction of the dataset is labeled. They are particularly useful for cerebrovascular segmentation, given that labeling a single volume requires several hours for an expert. In addition to the challenge posed by insufficient annotations, there are concerns regarding annotation consistency. The task of annotating the cerebrovascular tree is inherently ambiguous. Due to the discrete nature of images, the borders and extremities of vessels are often unclear. Consequently, annotations heavily rely on the expert subjectivity and on the underlying clinical objective. These discrepancies significantly increase the complexity of the segmentation task for the model and consequently impair the results. Consequently, it becomes imperative to provide clinicians with precise guidelines to improve the annotation process and construct more uniform datasets. In this article, we investigate the data dependency of deep learning methods within the context of imperfect data and semi-supervised learning, for cerebrovascular segmentation. Specifically, this study compares various state-of-the-art semi-supervised methods based on unsupervised regularization and evaluates their performance in diverse quantity and quality data scenarios. Based on these experiments, we provide guidelines for the annotation and training of cerebrovascular segmentation models.

*Keywords:* semi-supervised learning, cerebrovascular segmentation, annotation quality, benchmark

---

## 1. Introduction

In medical imaging, many pipelines used in clinical applications start by the segmentation of specific target structures. Consequently, the segmentation task has been widely investigated in the image processing community. In particular, deep-learning methods have successfully alleviated performance bottlenecks and yield state-of-the-art results [1]. Currently, most segmentation approaches developed in this context rely on the U-Net architecture comprising convolutional [2, 3, 4] and/or transformer-based [5, 6, 7] layers.

In the context of vascular segmentation [8, 9], cerebrovascular segmentation stands out as one of the most complex tasks. This complexity arises from the intricate nature of the brain vascular network, which comprises numerous tortuous vessels with a complex structure, featuring cycles and a tree-like topology. While conventional U-Net models have been applied to cerebrovascular segmentation [10, 11], recent advancements have introduced more sophisticated architectures [12, 13, 14, 15]. However, these approaches rely on supervised learning which requires a large quantity of labeled data, specifically the annotation of the whole brain vascular network across numerous 3D images.

Unfortunately, these annotations have to be produced and/or carefully reviewed by medical experts which is heavily time-consuming [16]. As a result, most cerebrovascular segmentation datasets contain a limited number of labeled training data, thereby limiting the performance of trained models.

In addition to the lack of labeled data, another significant challenge arises from the considerable variability in cerebrovascular labels. This variability comes from several factors. Firstly, the definition of a cerebrovascular label varies depending on the specific application of interest [17]. Some datasets concentrate solely on major vessels, such as the circle of Willis (TopCow [16]), while others include all arteries but exclude veins (Bullitt [18], IXI [19]). Despite the adoption of a global annotation policy, the limit regarding where to stop annotating can be ambiguous and vary from one dataset to another. For instance, when stating “all arteries are labeled”, some datasets may or may not include the ophthalmic arteries and the arteries supplying blood to the scalp.

Even though the same annotation policy is shared between two datasets, variations in the extent of labels may arise depending on the annotator. Firstly, due to the discrete nature of annotations, the status (inside or outside the vascular label) of voxels on the vessel borders becomes ambiguous and is subject to the annotator’s discretion. Secondly, a similar issue arises when labeling a vessel near the image resolution limit. The annotators may choose to stop labeling at different positions.

These numerous sources of vascular label variability create what is commonly referred to as concept shift [20, 21]. During training, the model tries to learn the concept of “vessel” based on the provided labels. However, this concept may undergo changes depending on the annotator responsible for its creation (inter-expert variability) or even the moment of the label creation (intra-expert variability). This concept shift is often even more important when using several datasets during the same training, as slight changes in the annotation policy

may lead to systematic bias in the concept of vessel. For example, the extent of the artery border as well as the arteries included in the label in the Bullitt dataset compared to the IXI dataset are significantly different (Fig. 1). This concept shift tends to reduce the performance of the segmentation model and the stability of the training. Additionally, it causes problems when evaluating a model on a dataset exhibiting a concept shift compared to the one used during training.

Secondly, extensively labeling cerebrovascular structures is a laborious task, making it prone to significant variations in the quality of annotations. Notably, common imperfections observed in cerebrovascular labels include missed vessels, disconnections or holes within a labeled vessel, and the inadvertent labeling of veins in datasets intended for arterial annotation (illustrated in Fig.1). Another common imperfection is the absence of spatial continuity in labels, especially when annotations are exclusively conducted in a single 2D plane, typically the axial one. These imperfections generate noisy labels, further complicating the training process of the segmentation model.

This problem of imperfect annotations is closely linked with the problem of annotations scarcity. When confronted with a limited quantity of labeled data, the absence of diversity in annotations can result in deep learning models exhibiting an overfitting behavior on the specific concept and/or noise present on the labels of the training dataset, potentially introducing the biases discussed above. Hence, there is a notable interest in examining these issues together.

As discussed before, the challenges of data scarcity, concept shift and noisy labels are more pronounced in cerebrovascular segmentation compared to other segmentation applications. Thus they constitute an important performance bottleneck for supervised segmentation methods. Yet, to the best of our knowledge, this has not been studied.

To address both the scarcity of cerebrovascular annotations and the challenges related to their quality and consistency, there has been a growing interest in semi-supervised methods. These methods tackle the problem of lack of annotation by leveraging unlabeled data, alongside labeled data, to increase the model performance. Three distinct strategies can be identified [22]: (1) generating pseudo-labels to train a supervised model [23, 24, 25]; (2) using unlabeled data to perform an unsupervised regularization; (3) employing unlabeled data to learn prior knowledge through self-supervised tasks [26, 27]. The most prevalent strategy among these is the unsupervised regularization approach, further categorized into four types [22]: (1) consistency learning [28, 29, 30, 31, 32, 33]; (2) co-training [34, 35]; (3) adversarial learning [36, 37, 38]; and (4) entropy minimization [39]. In this study, we only focus on the unsupervised regularization approaches, as they also address the challenge of data imperfection by limiting overfitting.

Semi-supervised with unsupervised regularization methods have been applied to cerebrovascular segmentation. For instance, [40] proposed a mean-teacher framework where the consistency loss was guided by a region connectivity model to introduce anatomical constraints. Also, [41] proposed to use, in addition to the segmentation model, a generative model that reconstructs the

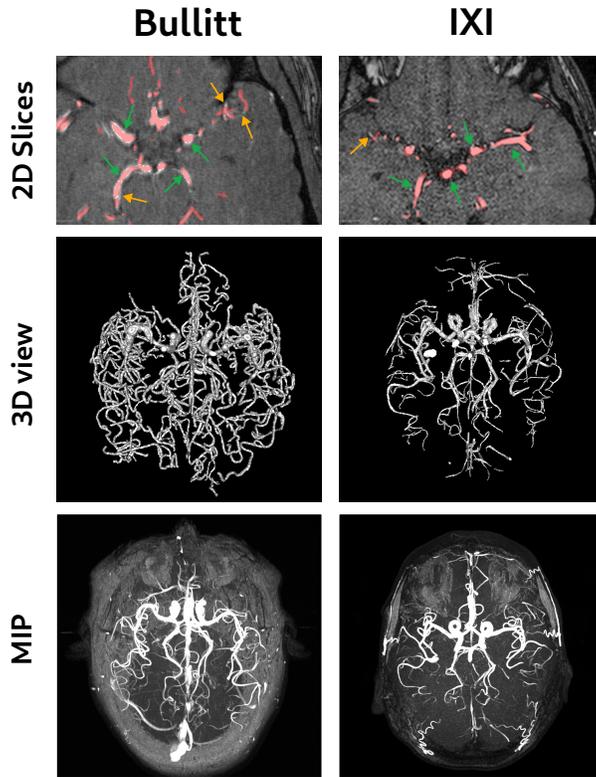


Figure 1: Illustration of noisy labels and concept shift problem arising on two cerebrovascular datasets : Bullitt [18] (left) and IXI [19] (right). First row: 2D slices with labels overlaid in light red, vessels with ambiguous boundaries indicated by green arrows and noisy labels (such as missing vessels or vessel disconnections), marked by orange arrows. Second row: 3D view showing the disparities in the extent of the labels for the same global annotation policy “labeling all cerebrovascular arteries”. Third row: maximum intensity projection (MIP) of the 3D volume.

input image from the segmentation outputs. A reconstruction loss was then computed to add another unsupervised regularization.

Employing semi-supervised cerebrovascular segmentation strategies seems appealing to cope with the lack of labeled data. Indeed, this approach not only boosts the model performance but also regularizes the training and limits overfitting. Nonetheless, it remains mandatory to include a sufficient amount of labeled data to enable the network to grasp the concept of a vessel. Moreover, this sufficient amount of labeled data highly depends on the quality of the annotations. Therefore, when encountering a new cerebrovascular segmentation challenge, finding a balance between the quantity of labeled data, their quality, and the chosen approach becomes paramount.

In this article, our goal is to formulate guidelines to assist the community in

finding this balance when addressing a cerebrovascular segmentation problem. In particular, we try to address the following questions:

- When is it more interesting to use a semi-supervised approach than a fully supervised approach for cerebrovascular segmentation?
- What type of semi-supervised approach should be used for cerebrovascular segmentation?
- Can semi-supervised learning methods effectively mitigate overfitting on the annotated data they use?
- Is the quality of annotation more crucial than quantity?
- What types of annotation imperfections most significantly degrade segmentation results and should be primarily avoided?

To do so, we identified five state-of-the-art semi-supervised segmentation approaches for medical images, along with a fully supervised baseline and compared them on two different cerebrovascular datasets. We specifically examine their behavior with respect to the number of available labeled data. We also study the sensibility of these approaches to the specific labeled samples that are used during training. We finally investigate the impact of several common imperfection types found in cerebrovascular labels: missing vessels, vessel border over- and under-segmentation.

The remainder of this article is organized as follows. Sect. 2 presents the datasets and semi-supervised methods used in this study. Sect. 3 details both the experimental setup and the outcomes of the conducted experiments. Finally, in Sect. 4, we discuss the results and conclusions of this study.

## 2. Materials and methods

In this section, we present the datasets as well as the employed semi-supervised learning approaches.

### 2.1. Datasets

We used the public dataset Bullitt [18] composed of 109 Time-of-Flight Magnetic Resonance Angiography (TOF-MRA) images of healthy patients acquired from the same center with a Siemens Allegra 3T system and 34 annotations. All volumes have the same resolution, with voxels of size  $0.51 \times 0.51 \times 0.80 \text{ mm}^3$  and a shape of  $448 \times 448 \times 128$ . We split this dataset into 94 volumes (including 19 annotated ones) for training, and 15 annotated volumes for test.

We also used the public dataset IXI [19] originally containing nearly 600 TOF-MRA images acquired from three different centers: Guy’s hospital, London, UK (316 patients), Hammersmith Hospital, London, UK (181 patients) and Institute of Psychiatry, London, UK (73 patients). Each of these patients is unique, and no patient underwent scans in different centers. For the three

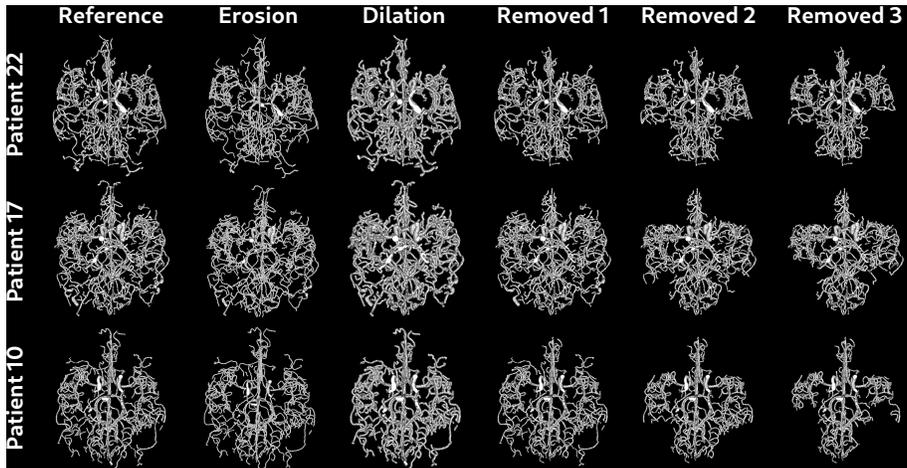


Figure 2: Visualization for three different patients of the different deteriorations carried out on the Bullitt dataset to investigate the impact of concept-shift in cerebrovascular segmentation.

centers, a limited number of annotations were available: fifteen, four and three respectively. We chose to restrict this study to one center in order to avoid the well-known performance drop linked to domain shift. Under these hypotheses, we used the 316 images from the Guy’s hospital center as it presented the highest number of annotations. All volumes were acquired at a consistent resolution with voxel dimensions of  $0.47 \times 0.47 \times 0.80 \text{ mm}^3$ , resulting in a size of  $512 \times 512 \times 100$  voxels. We randomly split the dataset into 311 images for training (301 without annotations and 10 with annotations), and 5 annotated images for testing.

For both datasets, the splitting between the training and test sets of the annotated volumes was performed randomly.

## 2.2. Metrics

In this study, our focus was restricted to two metrics. We first used the Dice score (*DSC*) [42], defined as:

$$DSC = \frac{2 \cdot tp}{2 \cdot tp + fp + fn} \quad (1)$$

where  $tp$ ,  $fp$ ,  $fn$  are the true positives, false positives and false negatives, respectively.

Additionally, we incorporated the cIDice score [43] that was built upon the Dice paradigm to mitigate bias towards large vessels, thus offering a more adapted evaluation for vessel segmentation, especially in the case of complete vascular networks that exhibit varying radius branches.

To compute the cIDice score, we first extract the centerlines  $C_{GT}$  and  $C_P$  from the ground truth segmentation mask  $S_{GT}$  and the predicted segmentation

$S_P$ , respectively. Then, we compute the ratio of  $C_P$  inside  $S_{GT}$ , called *topology precision* ( $Tp$ ), and the ratio of  $C_{GT}$  inside  $S_P$ , called *topology sensitivity* ( $Ts$ ):

$$Tp = \frac{|C_P \cap S_{GT}|}{C_P} \quad (2)$$

$$Ts = \frac{|C_{GT} \cap S_P|}{C_{GT}} \quad (3)$$

Then,  $clDice$  is defined as the harmonic mean of  $Tp$  and  $Ts$ :

$$clDice = \frac{2 \cdot Tp \cdot Ts}{Tp + Ts} \quad (4)$$

### 2.3. Methods

In this section, we present the five semi-supervised learning methods which were selected for comparison purposes. These methods all belong to the group of unsupervised regularization methods. They differ in their regularization strategy: consistency upon input noise [28, 30], adversarial consistency [36], task level consistency [29], and mutual consistency [31].

#### 2.3.1. Mean-teacher

Among the semi-supervised methods, the mean-teacher model [28] is a popular backbone (Fig. 3). In this framework, two different models are used, namely a teacher and a student. Only the student model is trained through back-propagation, while the teacher weights are updated by the exponential moving average of the student weights. The training procedure leverages the unlabeled data by performing an unsupervised regularization through a consistency loss. More precisely, the input data is disturbed with two different random Gaussian noises, leading to two slightly different inputs for the student and teacher models, respectively. Then, the consistency loss is used between the outputs of the student and teacher models, forcing these outputs to be consistent despite the perturbation. This regularization allows to leverage unlabeled data and improve the model performance. The model loss  $\mathcal{L}$  to optimize is:

$$\mathcal{L} = \mathcal{L}_s + \lambda_c \cdot \mathcal{L}_c \quad (5)$$

with  $\mathcal{L}_s$  the supervised segmentation loss computed on the labeled samples,  $\mathcal{L}_c$  the consistency loss computed on the unlabeled samples and  $\lambda_c \geq 0$  a parameter setting the weight of the consistency loss. For every method, the parameter  $\lambda_c$  evolves during training and follows a Gaussian warming up function.

#### 2.3.2. Uncertainty-aware mean-teacher

The uncertainty-aware mean-teacher [30] relies on the mean-teacher framework. It proposes to monitor the consistency regularization by using uncertainty. The uncertainty on the output of the teacher model is estimated with Monte-Carlo dropout. The model performs  $N$  forward passes with different dropouts, and the uncertainty is estimated by the entropy of these  $N$  predictions.

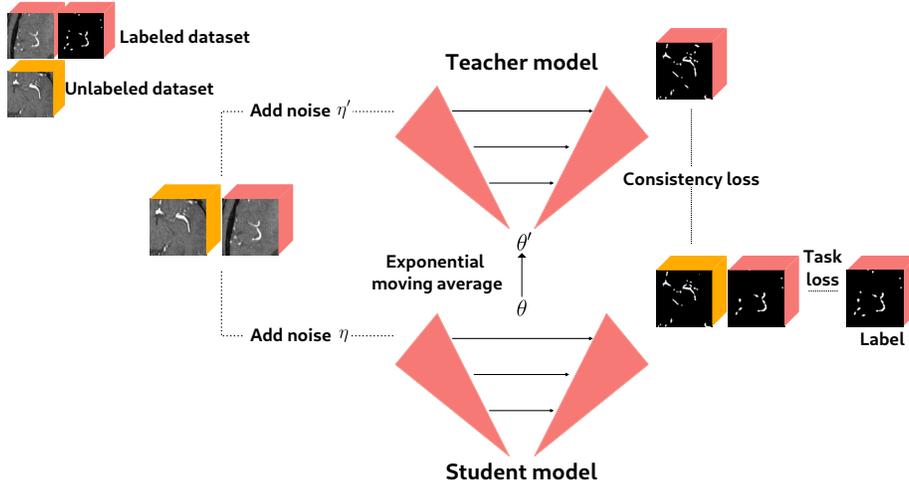


Figure 3: Illustration of the mean-teacher model leveraging both labeled and unlabeled data for cerebrovascular segmentation.

More formally, by setting  $P_n^c$  the probability volume of the  $c$ -th class of the  $n$ -th forward pass, the uncertainty volume is defined as:

$$U = - \sum_c \mathcal{U}_c \log(\mathcal{U}_c) \quad (6)$$

with:

$$\mathcal{U}_c = \frac{1}{N} \sum_n P_n^c \quad (7)$$

This uncertainty is used to monitor the consistency regularization. The computation of the consistency loss is restricted to voxels falling below a given threshold. Throughout training, this threshold is dynamically adjusted, ensuring that the consistency loss is computed only on voxels where the teacher’s prediction is reliable in the early stages of training, thereby enhancing regularization robustness. Conversely, as training progresses, the consistency loss is expanded to incorporate a broader set of ambiguous voxels, thereby capturing more complex information.

Let us consider  $\mathcal{P}_t$  the prediction of the teacher and  $\mathcal{P}_s$  the prediction of the student. We mask them with the uncertainty volume, as follows:

$$\mathcal{P}_t^{mask} = \mathbb{I}(U < \tau) \cdot \mathcal{P}_t \quad (8)$$

$$\mathcal{P}_s^{mask} = \mathbb{I}(U < \tau) \cdot \mathcal{P}_s \quad (9)$$

with  $\tau$  a dynamically adjusted threshold and  $\mathbb{I}$  the indicator function. The threshold  $\tau$  is defined by a Gaussian warming up function:

$$\tau(t) = \ln(2) \left( \frac{3}{4} + \frac{1}{4} \exp \left( -5 \left( 1 - \frac{t}{t_{max}} \right)^2 \right) \right)$$

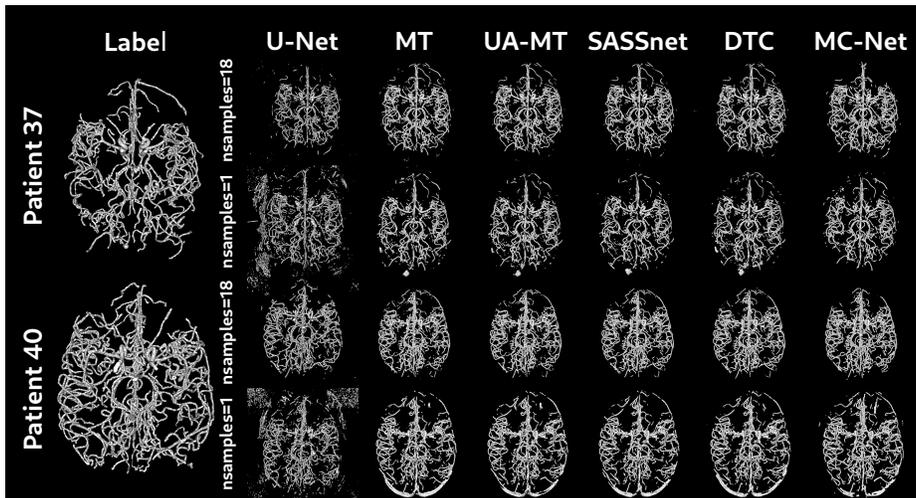


Figure 4: Visual results for Bullitt dataset from experiment 1. Segmentation results are presented for two patients, each method and two dataset compositions.

where  $t$  is the current training step and  $t_{max}$  the final training step.

The final model loss to optimize can be written as:

$$\mathcal{L} = \mathcal{L}_s + \lambda_c \cdot \mathcal{L}_c(\mathcal{P}_t^{mask}, \mathcal{P}_s^{mask}) \quad (10)$$

with  $\mathcal{L}_s$  the supervised segmentation loss computed on the labeled samples,  $\mathcal{L}_c$  the consistency loss and  $\lambda_c \geq 0$  a parameter setting the weight of the consistency loss.

### 2.3.3. Shape-aware semi-supervised network

The shape-aware semi-supervised (SASSnet) model [36] uses adversarial learning to enforce consistency of unlabeled and labeled data predictions. In this method, the regularization is not done on the segmentation output but on a signed distance map output to enforce geometric shape constraints.

Towards this goal, a multi-task architecture is built with a shared encoder and a decoder with two output heads: one for segmentation and one for signed distance map regression. The supervised loss  $\mathcal{L}_s$  is then the combination of the segmentation loss  $\mathcal{L}_{seg}$  and the mean squared error loss for the signed distance map  $\mathcal{L}_{sdm}$ :

$$\mathcal{L}_s = \mathcal{L}_{seg} + \alpha \cdot \mathcal{L}_{sdm} \quad (11)$$

with  $\alpha \geq 0$  a parameter balancing the two losses.

Then, a discriminator is designed to distinguish between the labeled and unlabeled datasets. More precisely, this discriminator receives the original image and the signed distance map as inputs and is required to classify whether the input originates from labeled data or unlabeled data. Optimizing this discriminator in an adversarial learning scheme helps to learn shape-aware features shared by both labeled and unlabeled datasets.

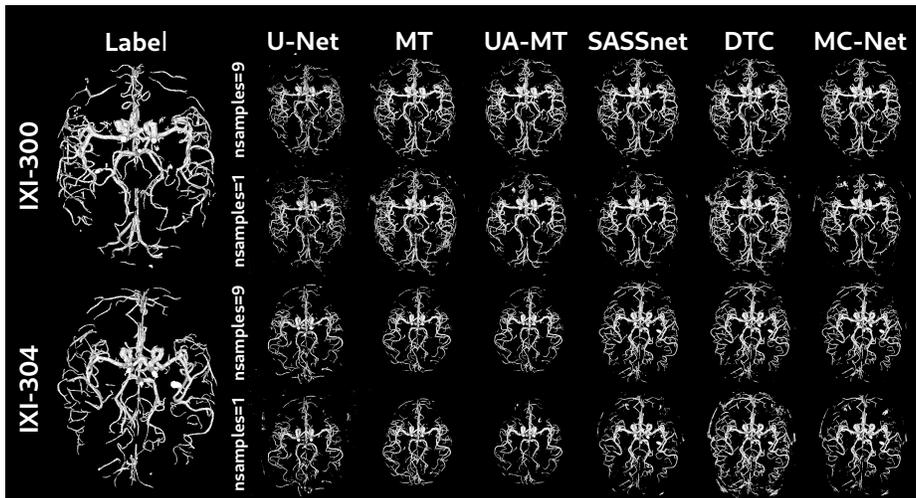


Figure 5: Visual results for IXI dataset from experiment 1. Segmentation results are presented for two patients, each method and two data dataset compositions.

The final training objective to optimize is written as:

$$\mathcal{L}(\theta, \zeta) = \mathcal{L}_s(\theta) + \lambda_c \cdot \mathcal{L}_{adv}(\theta, \zeta) \quad (12)$$

with  $\theta$  the parameters of the segmentation network,  $\zeta$  the parameter of the discriminator,  $\mathcal{L}_{adv}$  the adversarial loss and  $\lambda_c \geq 0$  a weight to control the importance of each loss. The optimization process is performed following:

$$\min_{\theta} \max_{\zeta} \mathcal{L}(\theta, \zeta) \quad (13)$$

#### 2.3.4. Dual-task consistency network

The dual-task consistency (DTC) network is based on task level consistency. Similarly to SASSnet [36], the framework consists of a voxel-wise segmentation head and a signed distance map regression head. The supervised loss is a combination of the segmentation loss  $\mathcal{L}_{seg}$  and the mean squared error loss for the signed distance map  $\mathcal{L}_{sdm}$ :

$$\mathcal{L}_s = \mathcal{L}_{seg} + \alpha \cdot \mathcal{L}_{sdm} \quad (14)$$

However, by contrast with SASSnet, the unsupervised regularization is performed by consistency regularization. Indeed, the signed distance map is transformed into a segmentation map, leading to a slightly different result than the segmentation head output. A consistency loss is finally used to enforce consistency between these two segmentation maps.

The mapping from the signed distance map to the segmentation map is approximated by the function  $\mathcal{T}(z_i)$  defined for a signed distance map value  $z_i$  of pixel  $i$  by:

$$\mathcal{T}(z_i) = (1 + e^{-k \cdot z_i})^{-1} \quad (15)$$

with  $k$  a parameter set to  $-1500$ . Then, the consistency loss  $\mathcal{L}_c$  is computed between the obtained segmentation map and the output of the segmentation head. The final model loss is defined as:

$$\mathcal{L}_{total} = \mathcal{L}_s + \lambda_c \cdot \mathcal{L}_c \quad (16)$$

with  $\mathcal{L}_s$  the segmentation loss computed on labeled data,  $\mathcal{L}_c$  the consistency loss and  $\lambda_c \geq 0$  a parameter setting the contribution of the consistency loss.

### 2.3.5. Mutual consistency network

The mutual consistency network (MC-Net) [31] consists of one shared encoder and two different decoders. The two decoders have a slightly different architecture. Since their outputs are different, the consistency loss enforces the consistency between them. More precisely, if we consider  $\mathcal{P}_A$  the prediction of the first decoder and  $\mathcal{P}_B$  the prediction of the second decoder,  $\mathcal{P}_A$  and  $\mathcal{P}_B$  are transformed into soft labels  $s\mathcal{P}_A$  and  $s\mathcal{P}_B$  through a sharpening function:

$$s\mathcal{P} = \frac{\mathcal{P}^{\frac{1}{T}}}{\mathcal{P}^{\frac{1}{T}} + (1 - \mathcal{P})^{\frac{1}{T}}} \quad (17)$$

with  $T$  a parameter to control the sharpening, set to 0.1.

Then  $s\mathcal{P}_A$  is used to supervise  $\mathcal{P}_B$  while  $s\mathcal{P}_B$  is used to supervise  $\mathcal{P}_A$ . This enforces the prediction consistency between the two decoders and promote low-entropy predictions. Finally, the loss to optimize is defined by:

$$\mathcal{L} = \mathcal{L}_{segA} + \mathcal{L}_{segB} + \lambda_c \cdot (\mathcal{L}_c(s\mathcal{P}_A, \mathcal{P}_B) + \mathcal{L}_c(s\mathcal{P}_B, \mathcal{P}_A)) \quad (18)$$

with  $\mathcal{L}_{segA}$  the supervised segmentation loss for the first decoder,  $\mathcal{L}_{segB}$  the supervised segmentation loss for the second decoder,  $\mathcal{L}_c$  the consistency loss and  $\lambda_c \geq 0$  a parameter setting the weight of the consistency loss.

## 3. Experiments and results

In this section, we explain in details the experiments carried out to investigate the semi-supervised learning approaches described above and to provide annotations guidelines for cerebrovascular segmentation.

### 3.1. Experimental setup

We performed three sets of experiments, aiming at addressing the questions raised in the introduction.

#### 3.1.1. Experiment 1

The purpose of the first experiment is to conduct a comparative analysis among various semi-supervised learning methods and assess their performance in comparison to a supervised baseline. More specifically, we aim to clarify in which context (i.e. with how many data available) semi-supervised learning brings a

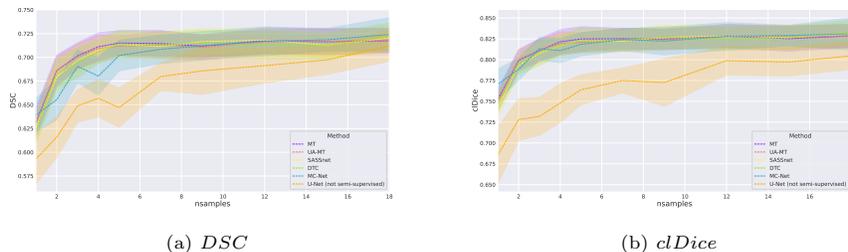


Figure 6: *DSC* (a) and *cLDice* (b) scores obtained on the test set for the Bullitt dataset, depending on the dataset composition. Each curve corresponds to a different method. The transparent bands represent the standard deviation.

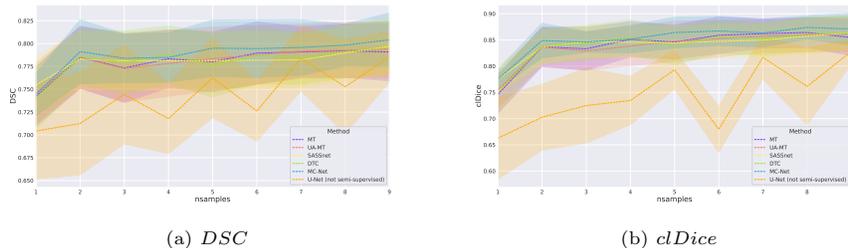


Figure 7: *DSC* (a) and *cLDice* (b) scores obtained on the test set for the IXI dataset, depending on the dataset composition. Each curve corresponds to a different method. The transparent bands represent the standard deviation.

significant improvement with respect to the supervised baseline. To this end, we evaluated all the methods on two datasets and with several dataset compositions (Tab. 3). Specifically, we increased the number of labeled data used for training while keeping the dataset size constant. For each data composition, one random training seed was selected, and one training was performed for each method.

### 3.1.2. Experiment 2

In a low-data regime, the risk of overfitting is very high and the performance of the model highly depends on the data used for training. We aim at investigating in which extent the semi-supervised learning framework succeeds to reduce both overfitting and dependency to the data used for training.

To achieve that goal, we compare a representative semi-supervised approach, the UA-MT method [30], to the supervised baseline with several different data seeds. We choose the UA-MT method among the semi-supervised methods as it is the most established method in the literature.

In this experiment, we selected the first five data compositions with the least number of labeled data (1 to 5). For each data composition and for each method (UA-MT and the baseline), we trained nine models, each with a different seed to select the labeled samples. The seeds were selected randomly but remained

| nsamples | U-Net      |               | MT         |               | UA-MT      |               | SASSnet    |               | DTC        |               | MC-Net     |               |
|----------|------------|---------------|------------|---------------|------------|---------------|------------|---------------|------------|---------------|------------|---------------|
|          | <i>DSC</i> | <i>clDice</i> |
| 1        | 0.593      | 0.685         | 0.632      | 0.753         | 0.631      | 0.750         | 0.626      | 0.742         | 0.624      | 0.747         | 0.639      | 0.771         |
|          | ± 0.025    | ± 0.033       | ± 0.015    | ± 0.012       | ± 0.012    | ± 0.010       | ± 0.013    | ± 0.013       | ± 0.014    | ± 0.011       | ± 0.017    | ± 0.018       |
| 2        | 0.616      | 0.728         | 0.686      | 0.799         | 0.686      | 0.800         | 0.679      | 0.793         | 0.683      | 0.795         | 0.655      | 0.789         |
|          | ± 0.021    | ± 0.024       | ± 0.016    | ± 0.012       | ± 0.014    | ± 0.012       | ± 0.016    | ± 0.011       | ± 0.015    | ± 0.012       | ± 0.021    | ± 0.013       |
| 3        | 0.649      | 0.732         | 0.702      | 0.811         | 0.701      | 0.812         | 0.697      | 0.807         | 0.694      | 0.809         | 0.691      | 0.813         |
|          | ± 0.017    | ± 0.022       | ± 0.013    | ± 0.014       | ± 0.013    | ± 0.012       | ± 0.015    | ± 0.013       | ± 0.015    | ± 0.015       | ± 0.017    | ± 0.013       |
| 4        | 0.657      | 0.748         | 0.711      | 0.822         | 0.709      | 0.819         | 0.707      | 0.817         | 0.706      | 0.817         | 0.680      | 0.811         |
|          | ± 0.019    | ± 0.023       | ± 0.014    | ± 0.014       | ± 0.014    | ± 0.014       | ± 0.012    | ± 0.013       | ± 0.012    | ± 0.011       | ± 0.019    | ± 0.014       |
| 5        | 0.647      | 0.764         | 0.715      | 0.825         | 0.712      | 0.821         | 0.714      | 0.823         | 0.714      | 0.823         | 0.702      | 0.818         |
|          | ± 0.020    | ± 0.017       | ± 0.012    | ± 0.014       | ± 0.010    | ± 0.011       | ± 0.012    | ± 0.012       | ± 0.014    | ± 0.015       | ± 0.016    | ± 0.013       |
| 7        | 0.680      | 0.775         | 0.715      | 0.825         | 0.713      | 0.824         | 0.712      | 0.822         | 0.713      | 0.823         | 0.708      | 0.824         |
|          | ± 0.014    | ± 0.014       | ± 0.013    | ± 0.012       | ± 0.014    | ± 0.013       | ± 0.014    | ± 0.012       | ± 0.013    | ± 0.013       | ± 0.014    | ± 0.016       |
| 9        | 0.686      | 0.773         | 0.712      | 0.825         | 0.711      | 0.824         | 0.717      | 0.827         | 0.715      | 0.827         | 0.712      | 0.822         |
|          | ± 0.024    | ± 0.028       | ± 0.011    | ± 0.014       | ± 0.013    | ± 0.012       | ± 0.014    | ± 0.015       | ± 0.013    | ± 0.013       | ± 0.014    | ± 0.016       |
| 12       | 0.691      | 0.799         | 0.718      | 0.828         | 0.716      | 0.828         | 0.718      | 0.828         | 0.713      | 0.825         | 0.717      | 0.828         |
|          | ± 0.018    | ± 0.017       | ± 0.013    | ± 0.014       | ± 0.012    | ± 0.013       | ± 0.013    | ± 0.012       | ± 0.014    | ± 0.014       | ± 0.015    | ± 0.016       |
| 15       | 0.697      | 0.797         | 0.716      | 0.825         | 0.717      | 0.827         | 0.713      | 0.827         | 0.712      | 0.826         | 0.718      | 0.829         |
|          | ± 0.015    | ± 0.016       | ± 0.013    | ± 0.014       | ± 0.013    | ± 0.013       | ± 0.014    | ± 0.014       | ± 0.013    | ± 0.014       | ± 0.012    | ± 0.012       |
| 18       | 0.712      | 0.805         | 0.717      | 0.829         | 0.718      | 0.828         | 0.720      | 0.831         | 0.722      | 0.833         | 0.724      | 0.831         |
|          | ± 0.016    | ± 0.016       | ± 0.013    | ± 0.014       | ± 0.012    | ± 0.013       | ± 0.013    | ± 0.013       | ± 0.014    | ± 0.014       | ± 0.017    | ± 0.019       |

Table 1: *DSC* and *clDice* for each method on Bullitt dataset, depending on the dataset composition.

| nsamples | U-Net      |               | MT         |               | UA-MT      |               | SASSnet    |               | DTC        |               | MC-Net     |               |
|----------|------------|---------------|------------|---------------|------------|---------------|------------|---------------|------------|---------------|------------|---------------|
|          | <i>DSC</i> | <i>clDice</i> |
| 1        | 0.704      | 0.663         | 0.742      | 0.746         | 0.746      | 0.755         | 0.755      | 0.775         | 0.741      | 0.749         | 0.744      | 0.775         |
|          | ± 0.047    | ± 0.070       | ± 0.029    | ± 0.032       | ± 0.029    | ± 0.029       | ± 0.028    | ± 0.022       | ± 0.031    | ± 0.032       | ± 0.021    | ± 0.016       |
| 2        | 0.712      | 0.703         | 0.785      | 0.837         | 0.785      | 0.836         | 0.783      | 0.836         | 0.785      | 0.838         | 0.791      | 0.849         |
|          | ± 0.050    | ± 0.056       | ± 0.030    | ± 0.033       | ± 0.030    | ± 0.034       | ± 0.028    | ± 0.029       | ± 0.027    | ± 0.030       | ± 0.031    | ± 0.030       |
| 3        | 0.744      | 0.725         | 0.773      | 0.833         | 0.773      | 0.829         | 0.781      | 0.847         | 0.779      | 0.844         | 0.784      | 0.846         |
|          | ± 0.048    | ± 0.064       | ± 0.034    | ± 0.036       | ± 0.033    | ± 0.032       | ± 0.028    | ± 0.026       | ± 0.031    | ± 0.030       | ± 0.024    | ± 0.017       |
| 4        | 0.718      | 0.734         | 0.783      | 0.851         | 0.782      | 0.849         | 0.782      | 0.850         | 0.788      | 0.855         | 0.785      | 0.852         |
|          | ± 0.034    | ± 0.041       | ± 0.028    | ± 0.030       | ± 0.028    | ± 0.032       | ± 0.028    | ± 0.028       | ± 0.028    | ± 0.030       | ± 0.024    | ± 0.023       |
| 5        | 0.762      | 0.793         | 0.780      | 0.846         | 0.783      | 0.848         | 0.781      | 0.844         | 0.776      | 0.838         | 0.795      | 0.864         |
|          | ± 0.039    | ± 0.033       | ± 0.029    | ± 0.029       | ± 0.031    | ± 0.034       | ± 0.029    | ± 0.029       | ± 0.031    | ± 0.032       | ± 0.028    | ± 0.026       |
| 6        | 0.726      | 0.680         | 0.790      | 0.859         | 0.782      | 0.851         | 0.788      | 0.857         | 0.788      | 0.857         | 0.794      | 0.863         |
|          | ± 0.029    | ± 0.039       | ± 0.030    | ± 0.031       | ± 0.031    | ± 0.030       | ± 0.029    | ± 0.030       | ± 0.028    | ± 0.029       | ± 0.028    | ± 0.024       |
| 7        | 0.784      | 0.817         | 0.791      | 0.862         | 0.792      | 0.857         | 0.782      | 0.850         | 0.784      | 0.851         | 0.796      | 0.863         |
|          | ± 0.032    | ± 0.037       | ± 0.025    | ± 0.025       | ± 0.026    | ± 0.028       | ± 0.026    | ± 0.028       | ± 0.025    | ± 0.027       | ± 0.027    | ± 0.020       |
| 8        | 0.753      | 0.761         | 0.792      | 0.864         | 0.790      | 0.861         | 0.794      | 0.864         | 0.792      | 0.861         | 0.798      | 0.874         |
|          | ± 0.046    | ± 0.065       | ± 0.026    | ± 0.025       | ± 0.026    | ± 0.027       | ± 0.026    | ± 0.026       | ± 0.026    | ± 0.029       | ± 0.021    | ± 0.020       |
| 9        | 0.785      | 0.828         | 0.791      | 0.855         | 0.794      | 0.861         | 0.797      | 0.864         | 0.796      | 0.863         | 0.804      | 0.871         |
|          | ± 0.025    | ± 0.022       | ± 0.028    | ± 0.031       | ± 0.028    | ± 0.030       | ± 0.028    | ± 0.031       | ± 0.026    | ± 0.028       | ± 0.026    | ± 0.027       |

Table 2: *DSC* and *clDice* for each method on Bullitt dataset, depending on the dataset composition.

consistent between the supervised and semi-supervised methods. Our goal is to quantify and compare the variability of both models, depending on the choice of the seed.

Table 3: Overview of the dataset sizes and the dataset compositions employed in Experiment 1.

|  | Bullitt  | IXI      |
|--|----------|----------|
| Total dataset size   | 109      | 316      |
| Training set size  | 94       | 311      |
| Test set size  | 15       | 5        |
| Labeled data in the training set                             | 19       | 10       |
| Various dataset compositions<br>(num labeled, num unlabeled) | (1, 93)  | (1, 310) |
|  | (2, 92)  | (2, 309) |
|  | (3, 91)  | (3, 308) |
|  | (4, 90)  | (4, 307) |
|  | (5, 89)  | (5, 306) |
|  | (7, 87)  | (6, 305) |
|  | (9, 85)  | (7, 304) |
|  | (12, 82) | (8, 303) |
|  | (15, 79) | (9, 302) |
|  | (18, 76) | –        |

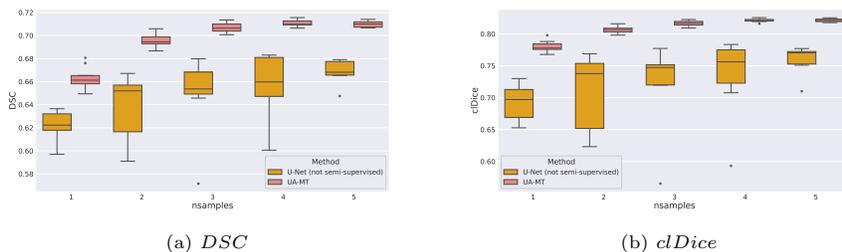


Figure 8: Boxplots of  $DSC$  (a) and  $cLDice$  (b) calculated on the different training seeds for Bullitt. Each color represents a different method (UA-MT and U-Net).

### 3.1.3. Experiment 3

In this experiment, our objective is to study the impact of both type and degree of label imperfection on the segmentation results. We aim to identify whether certain imperfections have a lower impact on segmentation outcomes. If so, the annotation process and review can be tailored to prioritize key factors that significantly affect results, optimizing both the annotation process and model performance. Ultimately, our goal is to formulate annotation guidelines for enhanced efficiency and model performance.

We identified three main categories of annotation imperfections: over-estimated vessel radius, under-estimated vessel radius, and missing small vessels (or part of vessels).

We simulated each type of imperfection as follows:

- Under-segmentation: all annotations are eroded with a ball of radius 1 voxel while ensuring that it does not introduce any vessel disconnection.

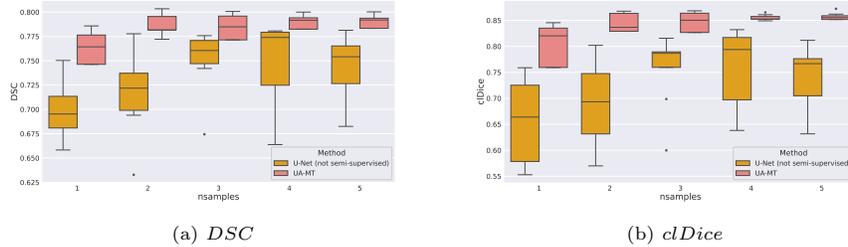


Figure 9: Boxplots of  $DSC$  (a) and  $cDice$  (b) calculated on the different training seeds for IXI. Each color represents a different method (UA-MT and U-Net).

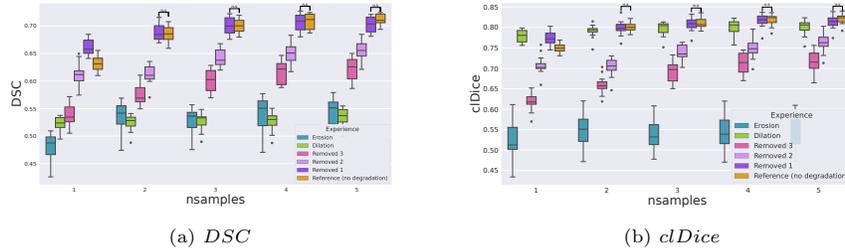


Figure 10: Boxplots of  $DSC$  (a) and  $cDice$  (b) on test set for each experiment, depending on the dataset composition. Each color correspond to a different experiment. Brackets indicate when the difference between results is not statistically significant (n.s.).

This scenario is referred to as “Erosion”.

- Over-segmentation: all annotations are dilated with a ball of radius 1. This scenario is referred to as “Dilation”.
- Missing vessels: we manually removed an increasing amount of vessels on the distal parts of the vascular network to create three levels of missing vessels. These scenarios are referred to as “Removed 1”, “Removed 2” and “Removed 3”.

These simulated imperfections were performed on 5 volumes. We then trained the UA-MT model [30] on these deteriorated datasets for increasing the number of available labeled samples and we tested it on the original dataset.

### 3.2. Implementation details

To conduct these experiments, we relied on the official implementation of the evaluated semi-supervised methods<sup>1</sup>. We adapted these implementations to our

<sup>1</sup><https://github.com/yulequan/UA-MT>  
<https://github.com/kleinzcy/SASSnet>

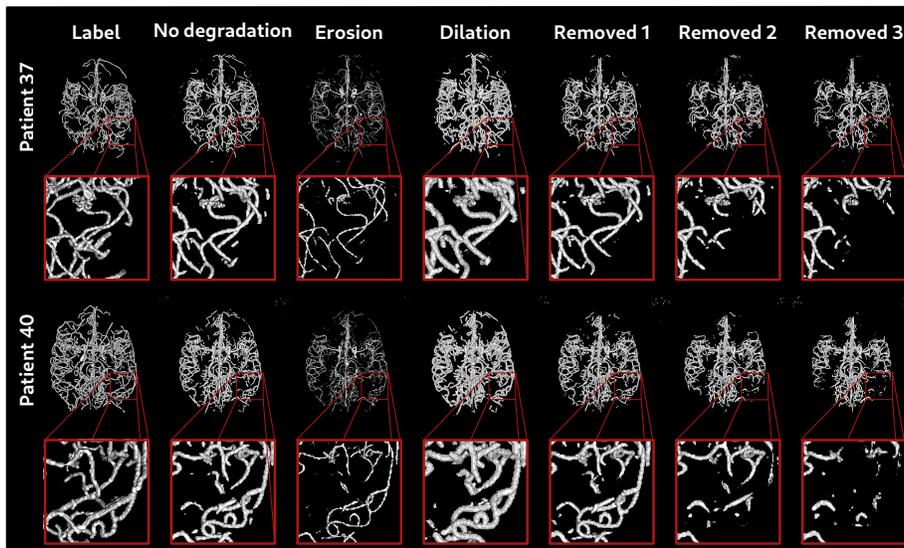


Figure 11: Visual results for two different patient from experiment 3. These segmentations are produced by the UA-MT model trained on the different damaged dataset (each column) with five training samples. Red boxes focus on some part of the cerebrovascular network to show better the differences between experiments.

datasets<sup>2</sup>. We also changed the VNet [44] backbone architecture by a standard U-Net to have a fair comparison with the state of the art U-Net [2] supervised baseline. All methods were trained with the same learning rate (0.01) and the same final weight  $\lambda_c$  for the consistency loss (0.01). Additionally, no data augmentation was performed, except for the Gaussian noise applied to introduce perturbations in the mean-teacher and uncertainty mean-teacher methods. All volumes were normalized using z-score. During training, patches were randomly extracted from the volume at each iteration to form a batch. The patch sizes were set to  $128 \times 128 \times 96$  for the IXI dataset and  $128 \times 128 \times 128$  for the Bullitt dataset. For the supervised loss on labeled data, we used the combination of soft Dice loss and cross-entropy whereas for the consistency loss, we used the soft Dice loss only.

### 3.3. Results

#### 3.3.1. Experiment 1: Methods comparison and impact of the quantity of annotated data

The results of these experiments are presented in Fig. 6 and Tab. 1 for the Bullitt dataset and Fig. 7 and Tab. 2 for the IXI dataset. Firstly, we can see that

<sup>1</sup><https://github.com/HiLab-git/DTC>  
<sup>2</sup><https://github.com/ycwu1997/MC-Net>  
<sup>2</sup><https://github.com/PierreRouge>

all semi-supervised methods perform better than the supervised U-Net baseline, regardless of the number of labeled samples in the training set. However, the degree of improvement diminishes with the addition of more labeled samples, as expected. For instance, the improvement in the *DSC* between UA-MT and U-Net when using 1 labeled sample is 0.04 but it decreases to 0.009 when using 9 labeled samples on the IXI dataset. Qualitative results confirm these findings (see Figs. 4 and 5). A statistical analysis reinforces these findings. Specifically, when examining the *DSC* values for the Bullitt dataset, all the p-values between U-Net and each semi-supervised method are lower than 0.001 across experiments involving one to seven samples. As the number of samples increases, the p-values rise above 0.05 when 18 samples are used. A similar trend is observed for the IXI dataset, though it is less pronounced due to fluctuations in U-Net results. In addition, we can observe that the segmentations generated by the supervised U-Net exhibit significant noise, as confirmed by the higher standard deviations, particularly when only one sample is used for training. This noise is typically attributed to an overfitting problem.

Another noteworthy result is the absence of a significant difference observed among the semi-supervised methods. Surprisingly, all methods exhibit similar performances irrespective of the number of labeled samples and for both datasets. This is supported by a statistical analysis. Indeed, when comparing the semi-supervised methods to each other, most p-values exceed 0.1, demonstrating that they yield similar results. This contrasts with findings in the literature, where more recent methods were reported to perform better on the LA dataset [45]. We assume that the higher complexity of the cerebrovascular structures and imperfections in the labeled data may cause these divergent conclusions. Conducting benchmarks on other types of datasets is guaranteed to draw more comprehensive and generalizable conclusions.

### 3.3.2. Experiment 2: Dependency to annotated data

The results of this experiment are presented in Fig. 8 and 9. For each data composition (i.e. number of labeled samples), we illustrated the segmentation results statistics over 9 trainings with different samples. This experiment confirms that the supervised baseline is less robust to the data variability than the semi-supervised approach, as indicated by the significantly lower standard deviation on both metrics, regardless of the number of labeled samples. This shows that the unsupervised regularization provided by the semi-supervised methods is effective to improve the generalization of the model.

Furthermore, while the addition of more labeled data contributes to regularization, intriguingly, the utilization of a substantial amount of unlabeled data appears to be more effective in this regard. Therefore, semi-supervised methods can serve to enhance performance to some extent, but their primary value lies in effectively regularizing the segmentation model.

### 3.3.3. Experiment 3: Impact of the annotation quality

The results of this experiment are presented in Fig. 10. The results show that training a model with annotations with a systematic under- or overesti-

mation of the vessel radius (dilation, removed 1/2/3) yields significantly worse results than with precise labels (reference). When only one labeled sample is used during training, the mean  $DSC$  on the test set drops from 0.63 (with the reference labels) to 0.52 and 0.48 with the dilated and eroded labels, respectively. This phenomenon persists regardless of the number of labeled samples in the training set. Moreover, the drop in  $DSC$  slightly increases with the addition of more labeled samples, indicating that the regularization typically observed when adding data does not compensate for the error in the labels. This was expected, as the labels present a systematic error in the experiment. Interestingly, the results obtained with the dilation experiment exhibit a  $clDice$  similar to that of the reference experiment. Indeed, the  $clDice$  is not a suitable metric to assess over-segmented results, as it compares the overlap of a centerline versus a segmentation. On the contrary, the results obtained with the erosion experiment exhibit a larger drop in  $clDice$  than in  $Dice$ . When removing the volumic bias of the  $Dice$  (i.e. where larger vessels weight more than thinner vessels), the results are even worse than expected. This indicates that many vessels are missing in the segmentation, which is confirmed by visual inspection (see Fig 11).

A compelling result from this experiment is that the segmentations from the removed 1/2/3 experiments are much better than that of dilation and erosion. This tends to demonstrate that when labeling vascular datasets, it is more important to accurately delineating the vessel borders than extensively annotating all vessels.

Indeed, we can see in Fig. 10 that when just some distal vessels are missed (experiment removed 1) the results are not significantly degraded. They are even improved for  $nsamples = 1$ , which again highlights the high sensibility of the results to the specific volume used for training (see experiment 2). In experiments removed 2 and removed 3, there is a significant degradation of the results, but much less important than for the other types of degradations. Moreover, contrary to the dilation and erosion experiments, the  $DSC$  increases regularly with the number of labeled samples. For  $nsamples = 2$ , the drop of  $DSC$  compared with the reference experiment is 0.08 and 0.12 for removed 2 and removed 3 respectively, and for  $nsamples = 5$ , the drop decreases to 0.05 and 0.09. The same observation can be seen for  $clDice$ . This shows a regularization effect when more labeled samples are added to the training set. The lack of some vessels in the labels introduces a noise in the annotations which is partially compensated when adding more partially annotated volumes. However, a performance gap still subsists because the missing vessels in the annotations introduce a concept-shift (i.e. small vessels located at the distal part of the vascular network are labeled “non vessels” in the training set but are labeled “vessels” in the test set) which is not compensated by the increasing number of data. This highlights the problem of concept shift that occurs when different experts have different understanding regarding the vessels to segment.

#### 4. Discussion and conclusion

In this study, we explored various semi-supervised methods across different data scenarios, aiming to discern the impact of both quantity and quality of annotations in the context of cerebrovascular segmentation. Our primary objective was to establish guidelines that optimize the annotation process and assist the community in making informed choices regarding model training. We addressed challenges such as annotation scarcity, noisy labels, and concept shift, particularly focusing on two common concept shift situations encountered in cerebrovascular segmentation: ambiguity at the vessel borders and uncertainty about which vessels should be included in the annotations.

Initially, our experiments indicated a significant advantage of using semi-supervised methods over a supervised model, especially when dealing with a low number of labeled samples. The extent of the benefit increased as the quantity of labeled samples decreased. However, we showed that the performance gap between semi-supervised methods and the supervised baseline diminished rapidly as the amount of annotated data increased. While this finding may temper expectations regarding the potential of semi-supervised methods to significantly enhance model performance, our experiments demonstrated their effectiveness in model regularization and reduction of overfitting. In the context of data scarcity, we believe that the most important added value of semi-supervised methods for cerebrovascular segmentation lies in these properties.

We shed light on the high impact of the concept shift in cerebrovascular segmentation. Such type of imperfection in the labeled data is very frequent in this application and appear when the annotation guidelines are not precise enough or when several experts perform the annotations. Contrary to the presence of noise in the labels that could be regularized by adding more labeled data, a concept shift problem, such as the under- or over-estimation of the vessel border, drastically decreases the performance of a model.

Based on our experiments, we can formulate several guidelines both regarding the annotation process of cerebrovascular images and the training of segmentation models. Firstly, when training a model using a non-in-house labeled dataset, it is very important to carefully analyze it to understand the chosen concept of a vessel during the annotation process. If the model exhibits poor generalization to a test dataset, the cause may be attributed to the concept shift between the training and test datasets. Secondly, when creating a new labeled dataset, it is of the utmost importance to define a precise policy to reduce the concept shift in the dataset. In particular, the following questions should be addressed: which vessels should be included in the annotation? When to stop labeling a vessel? What is the expected degree of precision for the vessel borders? Finally, the label quality is more important than their quantity. When there is a given time slot of expert availability, the primary focus should be on creating precise annotations, rather than annotating more samples with lower quality. In particular, precision in delineating vessel borders is more important than labeling extensively all the vessels up to the image resolution limit.

## Acknowledgements

This work was supported by the French *Agence Nationale de la Recherche* (grant ANR-20-CE45-0011) as part of the PreSPIN project. This work was also partly funded by France Life Imaging (grant ANR-11-INBS-0006). This work was granted access to the HPC resources of IDRIS under the allocation 2023-AD011013610 made by GENCI. We would like to thank Maria A. Zuluaga and Francesco Galati for graciously providing the IXI annotations.

## References

- [1] P.-H. Conze, G. Andrade-Miranda, V. K. Singh, V. Jaouen, D. Visvikis, Current and emerging trends in medical image segmentation with deep learning, *IEEE Transactions on Radiation and Plasma Medical Sciences* 7 (6) (2023) 545–569.
- [2] O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-Assisted Intervention (MICCAI), Proceedings, Part III, 2015*, pp. 234–241.
- [3] F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, K. H. Maier-Hein, nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation, *Nature Methods* 18 (2) (2021) 203–211.
- [4] R. Azad, E. K. Aghdam, A. Rauland, Y. Jia, A. H. Avval, A. Bozorgpour, S. Karimijafarbigloo, J. P. Cohen, E. Adeli, D. Merhof, Medical image segmentation review: The success of U-Net, *arXiv preprint arXiv:2211.14830* (2022).
- [5] J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, L. Lu, A. L. Yuille, Y. Zhou, TransUNet: Transformers make strong encoders for medical image segmentation, *arXiv preprint arXiv:2102.04306* (2021).
- [6] A. Hatamizadeh, V. Nath, Y. Tang, D. Yang, H. R. Roth, D. Xu, Swin UNETR: Swin transformers for semantic segmentation of brain tumors in MRI images, in: *International MICCAI Brainlesion Workshop, 2021*, pp. 272–284.
- [7] F. Shamsad, S. Khan, S. W. Zamir, M. H. Khan, M. Hayat, F. S. Khan, H. Fu, Transformers in medical imaging: A survey, *Medical Image Analysis* 88 (2023) 102802.
- [8] D. Lesage, E. D. Angelini, I. Bloch, G. Funka-Lea, A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes, *Medical Image Analysis* 13 (2009) 819–845.

- [9] S. Moccia, E. D. Momi, S. E. Hadji, L. S. Mattos, Blood vessel segmentation algorithms – Review of methods, datasets and evaluation metrics, *Computer Methods and Programs in Biomedicine* 158 (2018) 71–91.
- [10] P. Sanches, C. Meyer, V. Vigon, B. Naegel, Cerebrovascular network segmentation of MRA images with deep learning, in: *IEEE International Symposium on Biomedical Imaging (ISBI), Proceedings*, 2019, pp. 768–771.
- [11] M. Livne, J. Rieger, O. U. Aydin, A. A. Taha, E. M. Akay, T. Kossen, J. Sobesky, J. D. Kelleher, K. Hildebrand, D. Frey, V. I. Madai, A U-Net deep learning framework for high performance vessel segmentation in patients with cerebrovascular disease, *Frontiers in Neuroscience* 13 (2019) 97.
- [12] N. Valderrama, I. Pitsiorlas, L. Vargas, P. Arbelaez, M. A. Zuluaga, JOBVS: Joint brain-vessel segmentation in TOF-MRA images, in: *IEEE International Symposium on Biomedical Imaging (ISBI), Proceedings*, 2023.
- [13] F. Lin, Y. Xia, S. Song, N. Ravikumar, A. F. Frangi, High-throughput 3DRA segmentation of brain vasculature and aneurysms using deep learning, *Computer Methods and Programs in Biomedicine* 230 (2023) 107355.
- [14] L. Mou, Y. Zhao, H. Fu, Y. Liu, J. Cheng, Y. Zheng, P. Su, J. Yang, L. Chen, A. F. Frangi, M. Akiba, J. Liu, CS2-Net: Deep learning segmentation of curvilinear structures in medical imaging, *Medical Image Analysis* 67 (2021) 101874.
- [15] H. Zhang, L. Xia, R. Song, J. Yang, H. Hao, J. Liu, Y. Zhao, Cerebrovascular segmentation in MRA via reverse attention network, in: *Medical Image Computing and Computer Assisted Intervention (MICCAI), Proceedings, Part VI*, 2020, pp. 66–75.
- [16] K. Yang, F. Musio, Y. Ma, N. Juchler, J. C. Paetzold, R. Al-Maskari, L. Höher, H. B. Li, I. E. Hamamci, A. Sekuboyina, S. Shit, H. Huang, D. Waldmannstetter, F. Kofler, F. Navarro, M. Menten, I. Ezhov, D. Rueckert, I. Vos, Y. Ruigrok, B. Velthuis, H. Kuijff, J. Hämmerli, C. Wurster, P. Bijlenga, L. Westphal, J. Bisschop, E. Colombo, H. Baazaoui, A. Makmur, J. Hallinan, B. Wiestler, J. S. Kirschke, R. Wiest, E. Montagnon, L. Letourneau-Guillon, A. Galdran, F. Galati, D. Falcetta, M. A. Zuluaga, C. Lin, H. Zhao, Z. Zhang, S. Ra, J. Hwang, H. Park, J. Chen, M. Wodzinski, H. Müller, P. Shi, W. Liu, T. Ma, C. Yağın, R. E. Hamadache, J. Salvi, X. Llado, U. M. L.-T. Estrada, V. Abramova, L. Giancardo, A. Oliver, J. Liu, H. Huang, Y. Cui, Z. Lin, Y. Liu, S. Zhu, T. R. Patel, V. M. Tutino, M. Orouskhani, H. Wang, M. Mossa-Basha, C. Zhu, M. R. Rokuss, Y. Kirchhoff, N. Disch, J. Holzschuh, F. Isensee, K. Maier-Hein, Y. Sato, S. Hirsch, S. Wegener, B. Menze, Benchmarking the CoW with the Top-CoW challenge: Topology-aware anatomical segmentation of the circle of Willis for CTA and MRA, *arXiv preprint arXiv:2312.17670* (2023).

- [17] M. Ciecholewski, M. Kassjański, Computational methods for liver vessel segmentation in medical imaging: A review, *Sensors* 21 (6) (2021) 2027.
- [18] S. R. Aylward, E. Bullitt, Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction, *IEEE Transactions on Medical Imaging* 21 (2) (2002) 61–75.
- [19] <https://brain-development.org/ixi-dataset/>.
- [20] X. Liu, C. Yoo, F. Xing, H. Oh, G. El Fakhri, J.-W. Kang, J. Woo, Deep unsupervised domain adaptation: A review of recent advances and perspectives, *APSIPA Transactions on Signal and Information Processing* 11 (1) (2022).
- [21] W. M. Kouw, M. Loog, An introduction to domain adaptation and transfer learning, arXiv preprint arXiv:1812.11806 (2018).
- [22] R. Jiao, Y. Zhang, L. Ding, R. Cai, J. Zhang, Learning with limited annotations: A survey on deep semi-supervised learning for medical image segmentation, arXiv preprint arXiv:2207.14191 (2022).
- [23] Y. Shi, J. Zhang, T. Ling, J. Lu, Y. Zheng, Q. Yu, L. Qi, Y. Gao, Inconsistency-aware uncertainty estimation for semi-supervised medical image segmentation, *IEEE Transactions on Medical Imaging* 41 (3) (2021) 608–620.
- [24] X. Wang, Y. Yuan, D. Guo, X. Huang, Y. Cui, M. Xia, Z. Wang, C. Bai, S. Chen, SSA-Net: Spatial self-attention network for COVID-19 pneumonia infection segmentation with semi-supervised few-shot learning, *Medical Image Analysis* 79 (2022) 102459.
- [25] B. H. Thompson, G. Di Caterina, J. P. Voisey, Pseudo-label refinement using superpixels for semi-supervised brain tumour segmentation, in: *IEEE International Symposium on Biomedical Imaging (ISBI), Proceedings, 2022*.
- [26] C. You, Y. Zhou, R. Zhao, L. Staib, J. S. Duncan, SimCVD: Simple contrastive voxel-wise representation distillation for semi-supervised medical image segmentation, *IEEE Transactions on Medical Imaging* 41 (9) (2022) 2228–2237.
- [27] H. Zheng, L. Lin, H. Hu, Q. Zhang, Q. Chen, Y. Iwamoto, X. Han, Y.-W. Chen, R. Tong, J. Wu, Semi-supervised segmentation of liver using adversarial learning with deep atlas prior, in: *Medical Image Computing and Computer Assisted Intervention (MICCAI), Proceedings, Part VI, 2019*, pp. 148–156.
- [28] A. Tarvainen, H. Valpola, Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results, in: *International Conference on Neural Information Processing Systems (NeurIPS), Proceedings, 2017*, pp. 1195–1204.

- [29] X. Luo, J. Chen, T. Song, G. Wang, Semi-supervised medical image segmentation through dual-task consistency, in: AAAI Conference on Artificial Intelligence (AAAI), Proceedings, 2021, pp. 8801–8809.
- [30] L. Yu, S. Wang, X. Li, C.-W. Fu, P.-A. Heng, Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation, in: Medical Image Computing and Computer Assisted Intervention (MICCAI), Proceedings, Part II, 2019, pp. 605–613.
- [31] Y. Wu, M. Xu, Z. Ge, J. Cai, L. Zhang, Semi-supervised left atrium segmentation with mutual consistency training, in: Medical Image Computing and Computer Assisted Intervention (MICCAI), Proceedings, Part II, 2021, pp. 297–306.
- [32] T. Lei, D. Zhang, X. Du, X. Wang, Y. Wan, A. K. Nandi, Semi-supervised medical image segmentation using adversarial consistency learning and dynamic convolution network, *IEEE Transactions on Medical Imaging* 42 (2023) 1265–1277.
- [33] Y. Bai, D. Chen, Q. Li, W. Shen, Y. Wang, Bidirectional copy-paste for semi-supervised medical image segmentation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023, pp. 11514–11524.
- [34] X. Luo, M. Hu, T. Song, G. Wang, S. Zhang, Semi-supervised medical image segmentation via cross teaching between CNN and transformer, in: International Conference on Medical Imaging with Deep Learning (MIDL), Proceedings, 2022, pp. 820–833.
- [35] Y. Xia, F. Liu, D. Yang, J. Cai, L. Yu, Z. Zhu, D. Xu, A. Yuille, H. Roth, 3D semi-supervised learning with uncertainty-aware multi-view co-training, in: IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), Proceedings, 2020, pp. 3646–3655.
- [36] S. Li, C. Zhang, X. He, Shape-aware semi-supervised 3D semantic segmentation for medical images, in: Medical Image Computing and Computer Assisted Intervention (MICCAI), Proceedings, Part I, 2020, pp. 552–561.
- [37] Y. Zhang, L. Yang, J. Chen, M. Fredericksen, D. P. Hughes, D. Z. Chen, Deep adversarial networks for biomedical image segmentation utilizing unannotated images, in: Medical Image Computing and Computer Assisted Intervention (MICCAI), Proceedings, Part III, 2017, pp. 408–416.
- [38] J. Hou, X. Ding, J. D. Deng, Semi-supervised semantic segmentation of vessel images using leaking perturbations, in: IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), Proceedings, 2022, pp. 2625–2634.

- [39] W. Hang, W. Feng, S. Liang, L. Yu, Q. Wang, K.-S. Choi, J. Qin, Local and global structure-aware entropy regularized mean teacher model for 3D left atrium segmentation, in: *Medical Image Computing and Computer Assisted Intervention (MICCAI), Proceedings, Part I, 2020*, pp. 562–571.
- [40] L. Xie, Z. Chen, X. Sheng, Q. Zeng, J. Huang, C. Wen, L. Wen, G. Xie, Y. Feng, Semi-supervised region-connectivity-based cerebrovascular segmentation for time-of-flight magnetic resonance angiography image, *Computers in Biology and Medicine* 149 (2022) 105972.
- [41] C. Chen, K. Zhou, Z. Wang, R. Xiao, Generative consistency for semi-supervised cerebrovascular segmentation from TOF-MRA, *IEEE Transactions on Medical Imaging* 42 (2) (2022) 346–353.
- [42] L. R. Dice, Measures of the amount of ecologic association between species, *Ecology* 26 (1945) 297–302.
- [43] S. Shit, J. C. Paetzold, A. Sekuboyina, I. Ezhov, A. Unger, A. Zhylyka, J. P. W. Pluim, U. Bauer, B. H. Menze, cIDice—A novel topology-preserving loss function for tubular structure segmentation, in: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Proceedings, 2021*, pp. 16560–16569.
- [44] F. Milletari, N. Navab, S.-A. Ahmadi, V-Net: Fully convolutional neural networks for volumetric medical image segmentation, in: *International Conference on 3D vision (3DV), Proceedings, 2016*, pp. 565–571.
- [45] Z. Xiong, Q. Xia, Z. Hu, N. Huang, C. Bian, Y. Zheng, S. Vesal, N. Ravikumar, A. Maier, X. Yang, et al., A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging, *Medical image analysis* 67 (2021) 101832.