Acute and sub-acute stroke lesion segmentation from multimodal MRI

Albert Clèrigues*, Sergi Valverde, Jose Bernal, Jordi Freixenet, Arnau Oliver, Xavier Lladó

Institute of Computer Vision and Robotics, University of Girona, Spain

Abstract

Background and objective. Acute stroke lesion segmentation tasks are of great clinical interest as they can help doctors make better informed time-critical treatment decisions. Magnetic resonance imaging (MRI) is time demanding but can provide images that are considered the gold standard for diagnosis. Automated stroke lesion segmentation can provide with an estimate of the location and volume of the lesioned tissue, which can help in the clinical practice to better assess and evaluate the risks of each treatment.

Methods. We propose a deep learning methodology for acute and sub-acute stroke lesion segmentation using multimodal MR imaging. We pre-process the data to facilitate learning features based on the symmetry of brain hemispheres. The issue of class imbalance is tackled using small patches with a balanced training patch sampling strategy and a dynamically weighted loss function. Moreover, a combination of whole patch predictions, using a U-Net based CNN architecture, and high degree of overlapping patches reduces the need for additional post-processing.

Results. The proposed method is evaluated using two public datasets from the 2015 Ischemic Stroke Lesion Segmentation challenge (ISLES 2015). These involve the tasks of sub-acute stroke lesion segmentation (SISS) and acute stroke penumbra estimation (SPES) from multiple diffusion, perfusion and anatomical MRI modalities. The performance is compared against state-of-the-art methods with a blind online testing set evaluation on each of the challenges. At the time of submitting this manuscript, our approach is the first method in the online rankings for the SISS (DSC=0.59 ± 0.31) and SPES sub-tasks (DSC=0.84 ± 0.10). When compared with the rest of submitted strategies, we achieve top rank performance with a lower Hausdorff distance.

Conclusions. Better segmentation results are obtained by leveraging the anatomy and pathophysiology of acute stroke lesions and using a combined approach to minimize the effects of class imbalance. The same training procedure is used for both tasks, showing the proposed methodology can generalize well enough to deal with different unrelated tasks and imaging modalities without hyper-parameter tuning. In order to promote the reproducibility of our results, a public version of the proposed method has been released to the scientific community at `https://github.com/NIC-VICOROB/stroke-mri-segmentation`.

Keywords: Brain, MRI, ischemic stroke, automatic lesion segmentation, convolutional neural networks, deep learning

1. Introduction

Stroke is a medical condition by which an abnormal blood flow in the brain causes the death of cerebral tissue. Stroke is the third cause of morbidity worldwide, after myocardial infarction and cancer, and the most prevalent cause of acquired disability [1]. The affected tissue in the acute phase can be divided into three concentric regions depending on the potential for recovery, also referred as salvageability: core, penumbra and benign oligemia [2]. The core, located at the center, is formed by irreversibly damaged tissue from a fatally low blood supply. The penumbra, located around the core, represents tissue at risk but that can still be recovered if blood flow is quickly restored. Finally, the benign oligemia is the outer most ring whose vascularity has been altered but is not at risk of damage. Once the symptoms of stroke have been identified, a shorter time to treatment is highly correlated with a positive outcome [3]. Mechanical thrombectomy is a strongly recommended option for eligible patients [4]. However,

Corresponding author. A. Clèrigues, Ed. P-IV, Campus Montilivi, University of Girona, 17003 Girona (Spain). e-mail: albert.clerigues@udg.edu. Phone: +34 683645681; Fax: +34 972 418976.

this surgery is not free of risks. An overall complication rate of 15.3% was observed in a year long study [5]. In the treatment decision context, an estimate of the salvageable tissue can aid physicians take more informed treatment decisions.

The Ischemic Stroke Lesion Segmentation (ISLES) challenge started in 2015 to provide a platform for fair and direct comparison of automated methods. It included two sub-tasks, the sub-acute ischemic stroke lesion segmentation (SISS) and the acute stroke penumbra estimation (SPES). The following ISLES 2016 and 2017 editions changed its focus from lesion segmentation to chronic lesion outcome prediction from MRI. In the 2015 ISLES workshop results, the top three methods in the SPES sub-task all used Random Decision Forests (RDFs) [6] using hand-crafted features [7, 8, 9]. RDFs were typically used in methods for stroke lesion segmentation due to their excellent generalization properties, which make them well suited for difficult tasks with few training samples [10]. Recent advances on convolutional neural networks (CNNs) [11] have achieved superior results and are currently replacing RDFs in most state-of-the-art methods. In contrast with RDFs, CNNs enable the joint learning of optimal features and classification criteria at training time for the specific task. However, these kind of networks are still restricted by the architectural design, the amount and quality of available data and the training procedure. Recently, advances in regularization techniques and data imbalance handling allow for increased CNN generalization performance in brain lesion segmentation that rivals that of RDFs. The best method in the SISS sub-task of the 2015 ISLES workshop employed a deep learning strategy consisting of a dual path encoder network with a conditional random field (CRF) post-processing [12]. More recently, Zhang et al [13] achieved comparable results by using a similar CNN trained with a deep supervision technique and a multi-scale loss function. Despite the good results of these kind of networks, the U-Net architecture [14], an encoder-decoder network, is replacing other state-of-the-art architectures for stroke lesion segmentation. This is clearly seen in the submissions for the ISLES 2017 challenge, where 10 out of the 14 participating methods, including the top three, used CNNs based on the U-Net architecture [15].

In this work, we present a deep learning approach for acute and sub-acute stroke lesion segmentation from multimodal MRI images. We use a 3D asymmetric encoder-decoder network based on the U-Net architecture with global and local residual connections. Within our approach, the class imbalance issue is alleviated with the use of small patches with balanced training patch sampling strategies and a dynamically weighted loss function. Additionally, we pre-process the provided images to facilitate using the symmetry property of brain hemispheres. The methodology is evaluated by cross-validation with the training images and with a blind online testing set evaluation against other state-of-the-art methods. The proposed approach demonstrates state of the art performance by ranking first in the testing leaderboard of both challenges [16] without any dataset specific tuning.

2. Data

For evaluation of the proposed methodology we use the public datasets provided for the two sub-tasks of the 2015 ISLES challenge [17]. They both encompass stroke lesion segmentation tasks from MRI imaging but using different imaging modalities and acquisition time since onset.

2.1. SISS dataset

For the sub-acute ischemic stroke segmentation (SISS) sub-task, a dataset was provided with 28 training and 36 testing cases acquired in the first week after onset. For each case, 4 co-registered multimodal images were provided including anatomical (T1, T2, FLAIR) and diffusion (DWI) MRI. The images were acquired as 3D volumes of $230{\times}230{\times}153$ dimensions at $1{\times}1{\times}1$ mm spacing. All four MRI modalities were used for evaluation of the proposed approach. For the training images, the provided gold standard, the whole lesion extent, was manually segmented by an experienced medical doctor.

2.2. SPES dataset

The acute stroke penumbra estimation sub-task (SPES) included 30 training and 20 testing cases acquired in the first day after onset. For each case, 7 co-registered modalities were provided including anatomical (T1 contrast, T2), diffusion (DWI) and perfusion (CBF, CBV, TTP, Tmax) MRI. The images were acquired as 3D volumes of $96 \times 110 \times 71$ dimensions at $2 \times 2 \times 2$ mm spacing. All seven MRI modalities were used for evaluation of the proposed approach. For the training images, the gold standard segmentation, the penumbra label, was obtained as the mismatch between whole lesion extent and the core delineated in perfusion and diffusion images respectively.

2

3. Methodology

We propose a 3D patch based deep learning method using an asymmetrical residual CNN based on the U-Net architecture [14]. Within our approach, the class imbalance issue is addressed with a combination of techniques including the use of small patches ($24 \times 24 \times 16$) and a weighted loss function. We also regularize the training procedure with dropout [18], data augmentation and early stopping. For image segmentation, the use of whole patch predictions with a high degree of overlap minimizes the need for additional post-processing. In the following, we briefly describe the main components and implementation details of our methodology.

3.1. Data pre-processing

The given images are first pre-processed with a symmetric modality augmentation to allow learning of features based on the symmetry of brain hemispheres despite the small receptive field of the used patches. Explicit symmetry information was already shown to improve results for chronic stroke lesion segmentation [19]. In our case, instead of using one patch per hemisphere in a multi-path network we use a single joint patch with a single-path network. In practice, we augment the provided modality images with symmetric versions that swap the left and right hemispheres. We first flip one of the images along the mid-sagittal axis and then we apply FSL FLIRT [20] to perform a linear registration between the original and flipped image. Finally, the rest of modalities are registered using the same transformation. Figure 1 shows an example of the resulting symmetrically augmented modalities. These are then appended to the provided ones, effectively doubling the number of images for each patient. In this way, a single extracted patch will also include intensity information from the opposite hemisphere.

3.2. CNN architecture

The used architecture, illustrated in Figure 2, is a 3D asymmetrical encoder-decoder network based on the U-Net [14] architecture and its 3D extension, the 3D U-Net [21]. Additionally, we also use short and long residual connections as used by the 2D uResNet architecture [22] for chronic stroke in MRI. The asymmetry comes from the number of parameters found in the encoder and decoder branches, with 75% and 25% of the parameters respectively. We use residual blocks with two convolutional layers in the encoder and with a single convolutional for the decoder. It has been shown that the decoder's role is not as critical for segmentation, mainly upsampling the work of the encoder and fine-tuning the

details [23]. Additionally, instead of the more typical rectified linear unit (ReLU) [24] we use in our residual blocks a parametric version, the PReLU non-linearity [25], as suggested by Paszke et al [23]. We perform downsampling in each resolution step by concatenating the result of a max pooling operation and strided convolution as proposed by Szegedy et al [26]. This strategy avoids representational bottlenecks while keeping the number of parameters contained. Finally, upsampling in the decoder branch is performed with the use of transposed convolutions.

3.3. Class imbalance handling

The class imbalance issue is caused by the typically smaller extent of the lesion class as compared with the rest of healthy tissue class. If no deliberate action is taken, the training set will be composed mostly from examples of healthy tissue and few from the lesion. This would induce a biased learning that would harm the segmentation performance. To alleviate this issue, we use a combination of small patches with a balanced training patch sampling and a difficulty weighted loss function. The employed loss function, the Focal loss [27], is a dynamically weighted extension of the cross entropy loss defined as:

$$\text{FL}(p_t) = -\alpha_t(1 - p_t)^\gamma \log(p_t) \tag{1}$$

where p_t and α_t are the predicted probability and weight for class t respectively. This function is dynamically weighted inversely proportional to the prediction confidence, so the network learns less from confident classifications and more from misclassified examples. In this way, class imbalance is alleviated as the network stops learning from the larger amount of healthy examples while still learning from the less common lesioned tissue. We use the Focal loss default parameters as suggested by Lin et al [27], with scaling factor $\gamma = 2$ and class weights $\alpha_0 = 0.25$ and $\alpha_1 = 0.75$ for the healthy and lesion classes respectively.

The use of patches allows using a training sampling strategy that can undersample the healthy class and oversample the lesion for a more balanced class representation. The employed strategy is an extension of two recently proposed ones for brain lesions [28] and chronic stroke [22]. In practice, a goal number of patches to extract is set per patient, as we aim to have a balanced patch representation of each case. Then, 50% of the training patches are extracted centered on voxels corresponding to healthy tissue and the other 50% on lesion. These are sampled at regular spatial steps
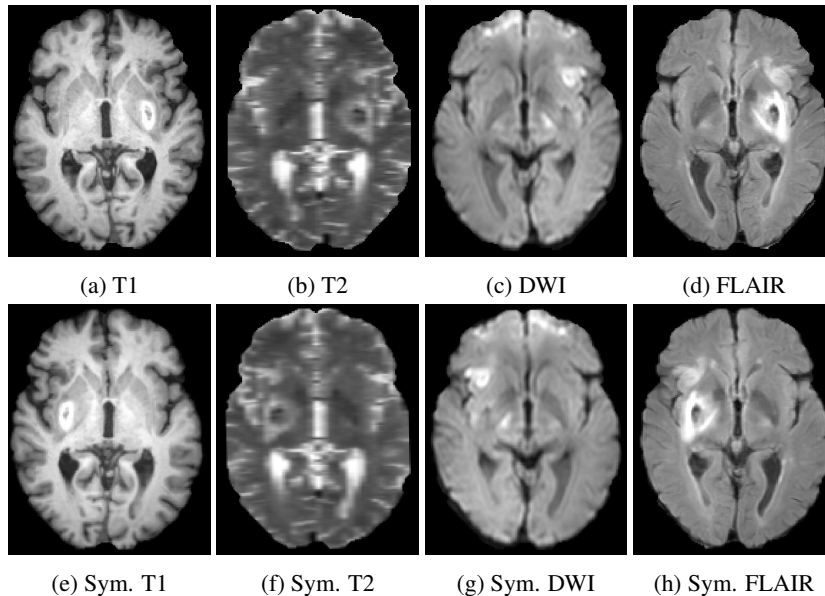
Figure 1: Provided and symmetrically augmented modalities from case 2 of the SISS training images.

to ensure that all parts of the brain are equally represented. The voxels sampled from the lesion class have a random offset added to increase representation of the region surrounding the lesion, the benign oligemia. As suggested by Guerrero et al [22], the offset is limited to half of the patch size to ensure the originally sampled voxel remains in the final extracted patch. For patients with smaller lesions, a combination of several patch extractions from the same lesion voxel and data augmentation is done to ensure the number is reached. The same sampled voxel will actually produce different patches since lesion voxels have a random offset applied. Finally, patches are extracted centered on these voxels. Additionally, for the lesion sampled patches, data augmentation is applied with five anatomically feasible operations including sagittal reflections and 90°, 180° and 270° axial rotations. A diagram summarizing the described strategy is depicted in Figure 3.

Despite the balancing effects of the Focal loss and training patch sampling, the segmentation performance is still reduced when bigger patch sizes are considered. Since there are much fewer lesion voxels than healthy ones, bigger patches tend to include more healthy class voxels and further worsen class imbalance in the training set. The employed patch size of $24 \times 24 \times 16$, determined empirically, offers the best compromise between receptive field and worsened imbalance for the considered datasets.

3.4. Network training

For training the randomly initialized network weights, we first extract patches to build the training and validation sets. As stated in Section 3.3, we use patches of size $24 \times 24 \times 16$ sampled with a balanced patch sampling strategy. During training, we use the Focal loss [27] along with the Adadelta optimizer [29], to avoid costly grid search of a learning rate, with a batch size of 16 patches. This optimizer requires no manual tuning of parameters and appears robust to noisy gradient information, different model architecture choices, various data modalities and selection of hyper-parameters. Moreover, to prevent overfitting we use the early stopping technique by monitoring the performance on a validation set at the end of each epoch. In this way, the training is interrupted when the monitored metric reaches a local minimum, which means no more generalizable knowledge is being learned from the training images. The sum of the L1 loss and error rate on the validation set is used as the monitored metric with a patience of 8 epochs.

3.5. Segmentation and post-processing

Once the network weights have been trained, to segment a new volume patches are first extracted from every part of the image and forward passed through the network. These are sampled uniformly with a regular extraction step of $4 \times 4 \times 1$ so that all parts of the brain are predicted. The resulting patch probabilities are then
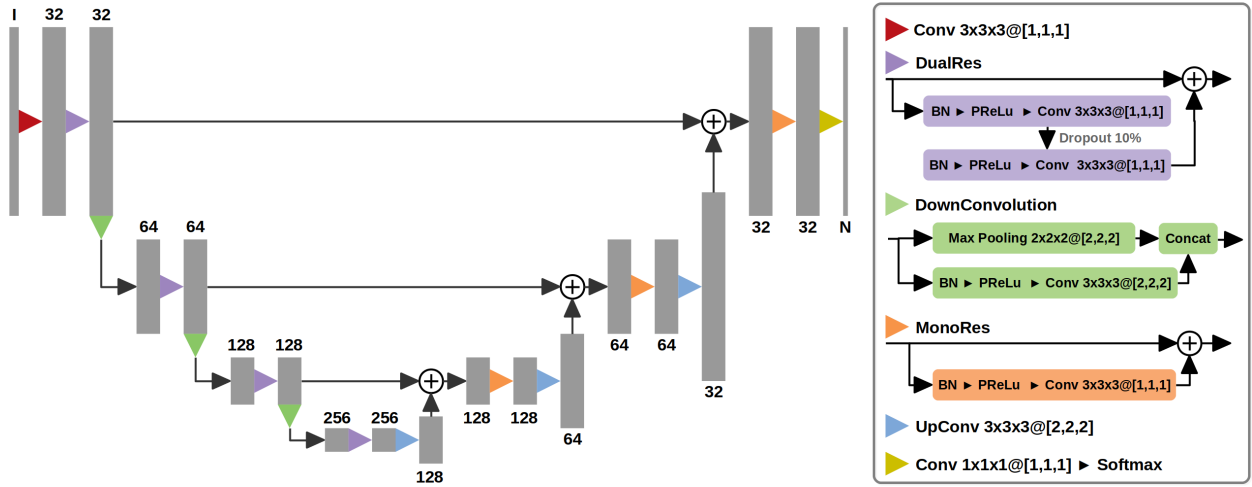
4

Figure 2: Employed U-Net based architecture using 3D convolutions, 4 resolution steps and 32 base filters. The architecture consists of an asymmetrical encoder-decoder network using long and short residual connections. For the convolutional layers, $K_x \times K_y \times K_z @[S_x, S_y, S_z]$ indicates the kernel and stride dimensions in each axis. The number of channels is indicated above or under each feature map. In the input and output feature maps, I and N denote the number of image modalities and segmentation classes respectively.
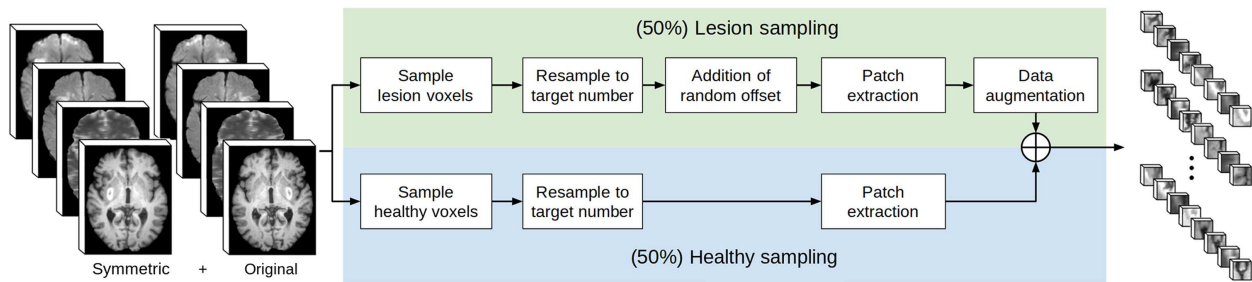

Figure 3: Diagram of the used training patch sampling strategy for acute stroke related tasks that considers the anatomy and pathophysiology of stroke lesions. The extracted patches are of size $24 \times 24 \times 16$ and include all input modalities.

combined in a common space preserving their original spatial location to produce the whole volume probability map. In our case, the combination is performed per voxel by averaging the class probabilities of the various patches. Furthermore, some degree of overlap between the extracted patches is used since the extraction step is smaller than the patch size. Therefore, the same voxel is labeled seen in different neighborhoods and the resulting class probabilities are averaged. This technique reduces the need for post-processing steps as it provides coherently spatial labels without block artifacts.

Finally, the probability maps are binarized by thresholding the lesion class probabilities and then performing a connected component filtering by lesion volume. The variable threshold T_h can compensate over/under confident networks while the minimum lesion size S_{min}, measured in number of voxels, takes advantage of lesion priors to minimise false positives. In practice, the probability maps are binarised using the same threshold

and minimum lesion size for each evaluation. These are found through grid search after all networks have been trained to offer the best compromise between the desired evaluation metrics.

3.6. Implementation details

The proposed method has been implemented with Python, using the Torch scientific computing framework [30]. All experiments have been run on a GNU/Linux machine running Ubuntu 18.04 with 64GB of RAM memory and an Intel® Core™ i7-7800X CPU. The network training and testing has been done with an NVIDIA TITAN X GPU (NVIDIA corp, United States) with 12GB G5X memory.

4. Evaluation and results

We perform a quantitative and qualitative evaluation with both a cross-validation experiment and a blind ex-

5

ternal evaluation using the challenge web platform. The metrics used in the quantitative evaluations will be the ones provided by the online platform. These include the Dice similarity coefficient (DSC) [31], sensitivity, positive predictive value (PPV) and Hausdorff distance (HD). The DSC measures the relative overlap of the segmentation with the ground truth and is used as a measure of segmentation performance. The sensitivity and PPV measure different properties relative to the lesion class segmentation. On the one hand, the sensitivity evaluates the percentage of gold standard lesion correctly labeled as such. On the other hand, the PPV measures the fraction of lesion class predictions that are correct. Finally, the Hausdorff distance can be intuitively seen as a measure of the *largest* border error between the segmentation and ground truth.

4.1. Cross-validation experiment with training images

The purpose of the cross-validation experiment is to quantitatively asses the main introduced improvements of the proposed methodology against a Baseline approach without them. For the Baseline approach, we use the proposed methodology without the class imbalance handling nor the data pre-processing step. Instead we use the crossentropy loss and training patch sampling as described in [28], using $24 \times 24 \times 16$ patches without any addition of a random offset. We then evaluate the effects of a Balanced approach that only uses the class imbalance handling described in Section 3.3, without performing symmetric modality augmentation. Finally, the Proposed approach also adds the data pre-processing step to implement the complete proposed methodology.

Each evaluation is performed in 4 folds, adjusting the number of cases per fold accordingly, with the same training procedure for both the SISS and SPES datasets. To build the patch training set for each fold, 10 000 patches per case are extracted from the training images summing approximately 260 000 patches in total. Once the networks from each fold have been trained and the probability maps generated for all training images, the post-processing parameters T_h and S_{min} are found through grid search to optimize the desired metrics across all folds. We consider the range of thresholds T_h from 0.1 to 0.9 and minimum lesion size S_{min} from 10 to 1000 voxels. More specifically, we choose the parameter combination that jointly maximizes the average DSC and HD, the two metrics used to determine the 2015 ISLES workshop results. In practice, a combined

score is computed as:

$$\text{Score} = \frac{\text{DSC} * \left(1 - \dfrac{\text{HD}}{\text{HD}_{\max}}\right)}{\text{DSC} + \left(1 - \dfrac{\text{HD}}{\text{HD}_{\max}}\right)} \qquad (2)$$

where $\text{HD}_{\max}$, set to 200 voxels, is used to normalize the HD metric to the range between 0 and 1.

4.1.1. SISS sub-task results

The evaluation metrics of the cross-validation experiment using the SISS dataset can be found in Table 1. With respect to the Baseline, the Balanced approach significantly improves the Hausdorff distance ($p < 0.01$) with marginal improvements in other metrics. When the symmetrically augmented modalities are further considered, the Proposed approach achieves significantly better DSC, PPV and HD ($p < 0.02$) as compared with the Baseline. However, despite the improvement in evaluation metrics, the Proposed approach needs a more restrictive minimum lesion size of 200 voxels to maximize the score as compared with the Baseline, which only filtered lesions smaller than 50 voxels.

Representative examples of the qualitative results from the proposed method can be found in Figure 4. Cases 9 and 15 represent the overall results of the proposed methodology, correctly detecting the lesions in most cases with an outline that approximates the provided gold standard. Among the observed limitations are inaccurate borders and over/under segmentation of certain regions. For instance, in case 5 the lesion was undersegmented due to a heterogeneous appearance of the gold standard lesion while in case 13 two false positive lesions are detected due to the previous existence of chronic stroke lesions with a similar appearance.

4.1.2. SPES sub-task results

The evaluation metrics of the cross-validation experiment using the SPES dataset can be found in Table 1. The class imbalance handling used in the Balanced approach significantly improves the sensitivity ($p < 0.01$) while providing marginal increase on the rest except the Hausdorff distance. When both improvements are simultaneously considered in the Proposed approach, it achieves a significantly better DSC and sensitivity ($p < 0.01$) than the Baseline. Additionally, the augmented modalities reduce the minimum lesion size S_{min} from 500 to a less restrictive 200 voxels.

Figure 5 shows qualitative results of four representative segmentation examples from the proposed method.

Table 1: Cross-validation experiment evaluation metrics on the SISS and SPES sub-tasks. The post-processing parameters T_h and S_{min} are found through grid search to maximize the score defined in Equation 1.

Approach	T_h	S_{min}	DSC	PPV	Sensitivity	HD
SISS sub-task						
Baseline	0.4	50	0.64 ± 0.22	0.69 ± 0.27	0.68 ± 0.21	43.7 ± 32.6
Balanced	0.4	200	0.67 ± 0.21	0.73 ± 0.22	0.69 ± 0.23	30.9 ± 28.9
Proposed	0.5	200	0.71 ± 0.19	0.78 ± 0.20	0.67 ± 0.22	29.5 ± 29.5
SPES sub-task						
Baseline	0.6	500	0.80 ± 0.17	0.82 ± 0.21	0.82 ± 0.19	11.1 ± 6.9
Balanced	0.4	500	0.82 ± 0.15	0.84 ± 0.14	0.85 ± 0.17	12.4 ± 7.6
Proposed	0.5	200	0.82 ± 0.16	0.85 ± 0.13	0.85 ± 0.17	11.2 ± 7.3

In general, the majority of the lesion is correctly segmented with minor border and small hole inaccuracies as seen in cases 11 and 26. Other less typical errors include under or oversegmentation of the lesion, as seen in case 2 where false positives are found on the upper part of the lesion. In the example of case 15, the lesion is undersegmented due to a confounding unusual appearance of some parts.

4.2. Blind challenge evaluation

To compare the proposed methodology against other state-of-the-art methods for acute stroke we submit our final approach for blind external evaluation in the ISLES 2015 challenge framework. The web platform used to hold the 2015 ISLES workshop [16] remains open for later submission and maintains an ongoing challenge leaderboard where the average testing set results are publicly displayed. Since the gold standard is hidden for the testing images, a fair and direct method comparison is possible. For evaluation in the challenge framework of ISLES 2015, we use the four networks trained for the Proposed approach during the cross-validation experiment, one from each fold, and average their outputs to produce a single testing patch prediction. The testing images are then segmented as described in Section 3.5 using the T_h and S_{min} set in the cross-validation experiment. In this way, the challenge results are produced with the same networks trained in the cross-validation experiment.

4.2.1. Challenge results

Tables 2 and 3 shows the top five entries as ranked by DSC of the ongoing testing leaderboard results for the SISS and SPES sub-tasks respectively. The proposed methodology achieves state-of-the-art performance in both sub-tasks, ranking first out of 74 entries in the SISS leaderboard and first out of 41 entries in the SPES leaderboard. As compared with the next best entries,

we achieve similar or higher DSC with a 12% and 28% lower Haussdorf distance in the SISS and SPES sub-tasks respectively. Additionally, in the SPES dataset we also obtain an 8% higher sensitivity.

5. Discussion

We have performed both qualitative and quantitative evaluations of the proposed methodology in two different tasks without any dataset specific tuning of training hyper-parameters. The methodology has been shown to perform equally well for the acute or sub-acute stages and with different combinations of MRI modalities. The results are improved with respect to the Baseline thanks to the combined approach to alleviate data imbalance and also through the explicit learning of features based on the brain symmetry. Additionally, the method is fast in inference, taking under 3 minutes to pre-process and predict a new image.

The proposed methodology demonstrates state-of-the-art performance ranking 1st by average DSC while having a smaller HD as compared with the next best method in both challenges. Moreover, we are the first U-Net based approach in the online testing leaderboard to outperform the best 2015 ISLES workshop entries. In the SISS sub-task, we obtain a similar DSC but with lower HD than the next best method. In contrast with the approach by Kamnitsas et al [12], we can avoid the use of the additional post-processing step with conditional random fields that needs several image dependent configurable parameters. In our case, the use of a U-Net based architecture that provides whole patch predictions allows performing highly overlapped segmentations without a large increase of inference time or introducing additional configurable parameters. In the SPES sub-task, we obtain a higher sensitivity with a lower HD as compared with the next best method by McKinley et al [8] that used a random decision forest classifier

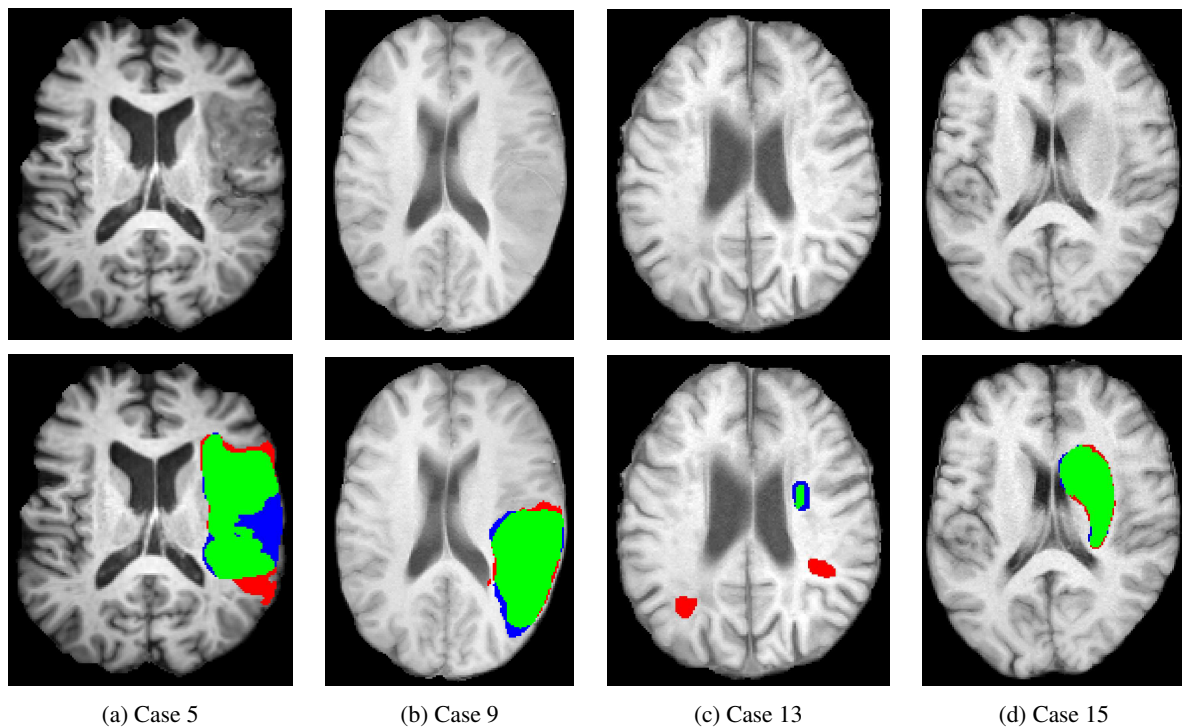|			
(a) Case 5	(b) Case 9	(c) Case 13	(d) Case 15

Figure 4: Output segmentation masks of representative cases from the training images of ISLES 2015 SISS dataset. On all images, true positives are denoted in green, false positives in red and false negatives in blue.

Table 2: Top 5 out of 74 entries of the ongoing SISS testing leaderboard [16] as ranked by average DSC.

Rank	User	DSC	PPV	Sensitivity	HD
1	**clera2 (ours)**	**0.59 ± 0.31**	**0.65 ± 0.35**	**0.60 ± 0.30**	**34.7 ± 28.9**
2	kamnk1 [12]	0.59 ± 0.31	0.68 ± 0.33	0.60 ± 0.27	39.6 ± 30.7
3	zhanr6 [13]	0.58 ± 0.31	0.60 ± 0.33	0.68 ± 0.24	38.9 ± 35.3
4	lianl1	0.57 ± 0.29	0.58 ± 0.30	0.64 ± 0.29	43.0 ± 30.5
5	saliz1	0.57 ± 0.31	0.54 ± 0.31	0.67 ± 0.29	41.1 ± 36.7

with several hand-crafted features over $3 \times 3 \times 3$ and $5 \times 5 \times 5$ neighbourhoods including local texture features, mean intensity, skewness, etc. By using a deep learning based method, the feature representation is learned at training time without having to rely on manually testing and finding the most appropriate ones for each specific task. Despite the good relative performance, the qualitative results show that the proposed methodology is still limited by inaccurate borders, missing lesion parts and other confounding factors. Furthermore, while the found minimum lesion size maximize the desired metrics along all training images they might still filter out some small lesions at testing time.

6. Conclusions

In this work, we have presented a methodology that achieves state-of-the-art performance in two different stroke lesion segmentation tasks. To the best of our knowledge, the proposed methodology is the first to obtain competitive results in both the ISLES 2015 SISS and SPES sub-tasks with the same approach. We have achieved these results by doing both regularization of the training procedure and providing additional meaningful information for lesion segmentation. Useful features using the brain symmetry could not be learned as the employed patch size is too small to include both hemispheres. The proposed symmetric modality augmentation facilitates using the similarity between hemispheres to improve lesion localization without using larger patches that would worsen class imbalance. The

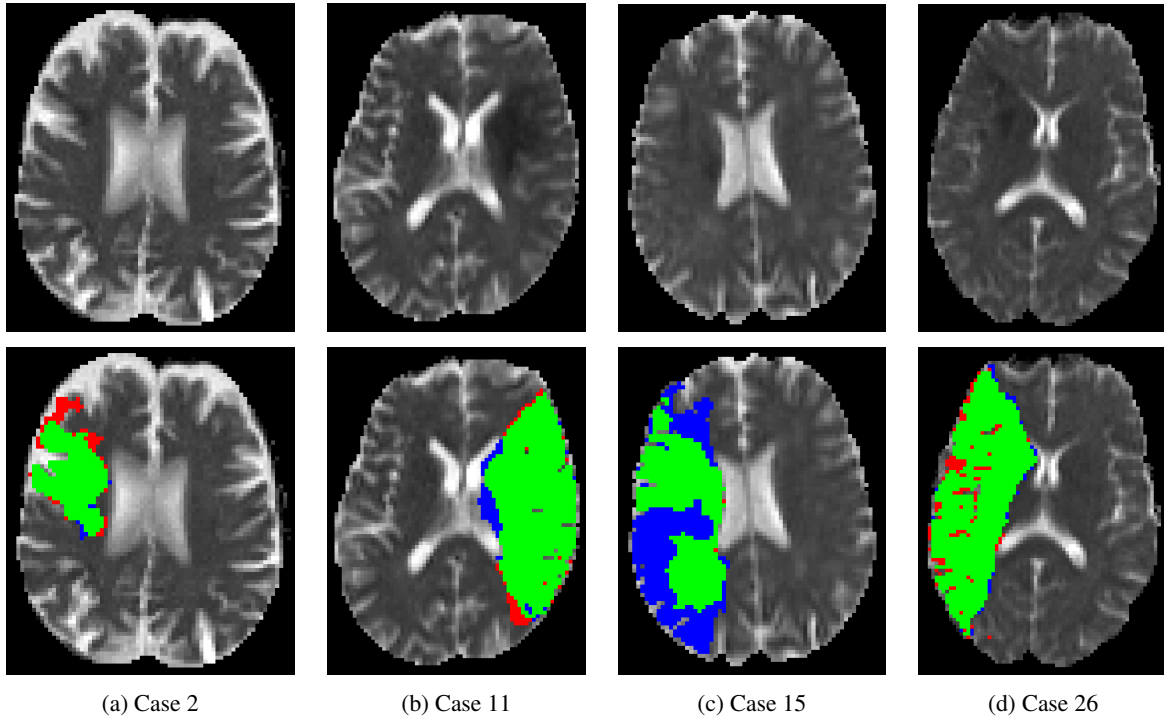|			
(a) Case 2	(b) Case 11	(c) Case 15	(d) Case 26

Figure 5: Output segmentation masks of representative cases from the training images of ISLES 2015 SPES dataset. On all images, true positives are denoted in green, false positives in red and false negatives in blue.

Table 3: Top 5 entries out of 41 of the ongoing SPES testing leaderboard [16] as ranked by average DSC.

Rank	User	DSC	PPV	Sensitivity	HD
1	**clera2 (ours)**	**0.84 ± 0.10**	**0.82 ± 0.15**	**0.89 ± 0.06**	**20.7 ± 13.9**
2	mckir1 [8]	0.82 ± 0.08	0.83 ± 0.10	0.82 ± 0.14	29.0 ± 16.3
3	cheng5	0.81 ± 0.11	0.81 ± 0.12	0.81 ± 0.14	22.7 ± 12.6
4	maieo1 [7]	0.81 ± 0.09	0.84 ± 0.08	0.80 ± 0.14	23.6 ± 13.0
5	ghosp1	0.80 ± 0.11	0.80 ± 0.15	0.83 ± 0.11	57.1 ± 25.4

use of additional informative modalities can be generalized with different augmentation techniques to facilitate the learning of more discriminative and meaningful features for the task at hand. Moreover, we have shown the big influence class imbalance can have in reducing distant outliers and false positives that provide a lower Hausdorff distance at testing time. By using a combined approach we achieve a less biased segmentation with a better balance between sensitivity and specificity. In the clinical setting, deep learning based methods can additionally benefit from related techniques such as transfer learning to learn an unrelated task with a lack of training examples [32] or to perform domain adaptation with few images [33]. The proposed methodology is made publicly available for the scientific community [34].

References

[1] J. Redon, M. H. Olsen, R. S. Cooper, O. Zurriaga, M. A. Martinez-Beneito, S. Laurent, R. Cifkova, A. Coca, and G. Mancia. Stroke mortality and trends from 1990 to 2006 in 39 countries from Europe and Central Asia: implications for control of

high blood pressure. *European Heart Journal*, 32(11):1424–1431, 2011. doi: 10.1093/eurheartj/ehr045.

[2] Islem Rekik, Stéphanie Allassonnière, Trevor K. Carpenter, and Joanna M. Wardlaw. Medical image analysis methods in MR/CT-imaged acute-subacute ischemic stroke lesion: Segmentation, prediction and insights into dynamic evolution simulation models. A critical appraisal. *NeuroImage: Clinical*, 1(1): 164–178, 2012. doi: 10.1016/J.NICL.2012.10.003.

[3] Sunil A. Sheth, Reza Jahan, Jan Gralla, Vitor M. Pereira, Raul G. Nogueira, Elad I. Levy, Osama O. Zaidat, and Jeffrey L. Saver. Time to endovascular reperfusion and degree of disability in acute stroke. *Annals of Neurology*, 78(4):584–593, 2015. doi: 10.1002/ana.24474.

[4] Bruce CV Campbell, Peter J Mitchell, Leonid Churilov, Mahsa Keshtkaran, Keun-Sik Hong, Timothy J Kleinig, Helen M Dewey, Nawaf Yassi, Bernard Yan, and Richard J Dowling. Endovascular thrombectomy for ischemic stroke increases disability-free survival, quality of life, and life expectancy and reduces cost. *Frontiers in Neurology*, 8:657, 2017. doi: 10.3389/fneur.2017.00657.

[5] Rakeshsingh K Singh, Vishal Annaji Chafale, Rakesh Shyam Lalla, Keyurkumar Chandrakantbhai Panchal, Anil Pandurang Karapurkar, Satish Vasant Khadilkar, Pawan K Ojha, Yogesh Godge, Rakesh K Singh, and Rajesh Benny. Acute Ischemic Stroke Treatment Using Mechanical Thrombectomy: A Study of 137 Patients. *Annals of Indian Academy of Neurology*, 20(3): 211–216, 2017. doi: 10.4103/aian.AIAN_158_17.

[6] Tin Kam Ho. Random decision forests. In *Proceedings of 3rd International Conference on Document Analysis and Recognition*, volume 1, pages 278–282. IEEE Comput. Soc. Press, 1995. ISBN 0-8186-7128-9. doi: 10.1109/ICDAR.1995.598994.

[7] Oskar Maier, Matthias Wilms, and Heinz Handels. Random forests for acute stroke penumbra estimation. In *Proceedings of ISLES 2015 challenge*, pages 77–80, 2015.

[8] Richard Mckinley, Levin Häni, Roland Wiest, and Mauricio Reyes. Segmenting the ischemic penumbra: a spatial Random Forest approach with automatic threshold finding. In *Proceedings of ISLES 2015 challenge*, pages 69–73, 2015.

[9] David Robben, Daan Christiaens, Janaki Raman Rangarajan, Jaap Gelderblom, Philip Joris, Frederik Maes, Paul Suetens, Oskar Maier, Mauricio Reyes, David Robben, Daan Christiaens, Janaki Raman Rangarajan, Jaap Gelderblom, Philip Joris, Frederik Maes, and Paul Suetens. ISLES Challenge 2015: A voxelwise, cascaded classification approach to stroke lesion segmentation. In *Proceedings of ISLES 2015 challenge*, pages 27–31. Springer, Cham, 2015. doi: 10.1007/978-3-319-30858-6_22.

[10] Oskar Maier, Christoph Schröder, Nils Daniel Forkert, Thomas Martinetz, and Heinz Handels. Classifiers for Ischemic Stroke Lesion Segmentation: A Comparison Study. *PLOS ONE*, 10 (12):e0145118, 2015. doi: 10.1371/journal.pone.0145118.

[11] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation*, 1(4): 541–551, 1989. doi: 10.1162/neco.1989.1.4.541.

[12] Konstantinos Kamnitsas, Liang Chen, Christian Ledig, Daniel Rueckert, and Ben Glocker. Multi-Scale 3D Convolutional Neural Networks for Lesion Segmentation in Brain MRI. In *Proceedings of ISLES 2015 challenge*, pages 13–17, 2015.

[13] Rongzhao Zhang, Lei Zhao, Wutao Lou, Jill M. Abrigo, Vincent C. T. Mok, Winnie C. W. Chu, Defeng Wang, and Lin Shi. Automatic Segmentation of Acute Ischemic Stroke From DWI Using 3-D Fully Convolutional DenseNets. *IEEE Transactions on Medical Imaging*, 37(9):2149–2160, 2018. doi: 10.1109/TMI.2018.2821244.

[14] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[15] Stefan Winzeck, Arsany Hakim, Richard McKinley, Jose A. Pinto, Victor Alves, Carlos Silva, Maxim Pisov, Egor Krivov, Mikhail Belyaev, Miguel Monteiro, Arlindo Oliveira, Youngwon Choi, Myunghee Cho Paik, Yongchan Kwon, Hanbyul Lee, Beom Joon Kim, Joong-Ho Won, Mobarakol Islam, Hongliang Ren, David Robben, Paul Suetens, Enhao Gong, Yilin Niu, Junshen Xu, John M. Pauly, Christian Lucas, Mattias P. Heinrich, Luis C. Rivera, Laura S. Castillo, Laura A. Daza, Andrew L. Beers, Pablo Arbelaezs, Oskar Maier, Ken Chang, James M. Brown, Jayashree Kalpathy-Cramer, Greg Zaharchuk, Roland Wiest, and Mauricio Reyes. ISLES 2016 and 2017-Benchmarking Ischemic Stroke Lesion Outcome Prediction Based on Multispectral MRI. *Frontiers in Neurology*, 9, 2018. doi: 10.3389/fneur.2018.00679.

[16] ISLES 2015 - SICAS Medical Image Repository, Leaderboard: SISS/SPES Testing., 2019 (accessed: 3 April 2019). URL https://www.smir.ch/ISLES/Start2015.

[17] Oskar Maier, Bjoern H. Menze, Janina von der Gablentz, Levin Häni, Mattias P. Heinrich, Matthias Liebrand, Stefan Winzeck, Abdul Basit, Paul Bentley, Liang Chen, Daan Christiaens, Francis Dutil, Karl Egger, Chaolu Feng, Ben Glocker, Michael Götz, Tom Haeck, Hanna-Leena Halme, Mohammad Havaei, Khan M. Iftekharuddin, Pierre-Marc Jodoin, Konstantinos Kamnitsas, Elias Kellner, Antti Korvenoja, Hugo Larochelle, Christian Ledig, Jia-Hong Lee, Frederik Maes, Qaiser Mahmood, Klaus H. Maier-Hein, Richard McKinley, John Muschelli, Chris Pal, Linmin Pei, Janaki Raman Rangarajan, Syed M.S. Reza, David Robben, Daniel Rueckert, Eero Salli, Paul Suetens, Ching-Wei Wang, Matthias Wilms, Jan S. Kirschke, Ulrike M. Krämer, Thomas F. Münte, Peter Schramm, Roland Wiest, Heinz Handels, and Mauricio Reyes. ISLES 2015 - A public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRI. *Medical Image Analysis*, 35:250–269, 2017. doi: 10.1016/J.MEDIA.2016.07.009.

[18] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014. doi: 10.1214/12-AOS1000.

[19] Yanran Wang, Aggelos K. Katsaggelos, Xue Wang, and Todd B Parrish. A deep symmetry convnet for stroke lesion segmentation. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 111–115. IEEE, 2016. ISBN 978-1-4673-9961-6. doi: 10.1109/ICIP.2016.7532329.

[20] Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–41, 2002.

[21] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 424–432. Springer, Cham, 2016. doi: 10.1007/978-3-319-46723-8_49.

[22] R. Guerrero, C. Qin, O. Oktay, C. Bowles, L. Chen, R. Joules, R. Wolz, M.C. Valdés-Hernández, D.A. Dickie, J. Wardlaw, and D. Rueckert. White matter hyperintensity and stroke lesion segmentation and differentiation using convolutional neural networks. *NeuroImage: Clinical*, 17:918–934, 2018. doi: 10.1016/J.NICL.2017.12.022.

[23] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. ENet: A Deep Neural Network Architec-

ture for Real-Time Semantic Segmentation. *arXiv preprint arXiv:1606.02147*, 2016.

[24] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.

[25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

[26] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2015. doi: 10.1109/CVPR.2016.308.

[27] Tsung-Yi Lin, Priyal Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. doi: 10.1109/TPAMI.2018.2858826.

[28] Konstantinos Kamnitsas, Christian Ledig, Virginia F.J. Newcombe, Joanna P. Simpson, Andrew D. Kane, David K. Menon, Daniel Rueckert, and Ben Glocker. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Medical Image Analysis*, 36:61–78, 2017. doi: 10.1016/J.MEDIA.2016.10.004.

[29] Matthew D. Zeiler. ADADELTA: An Adaptive Learning Rate Method. *arXiv preprint arXiv:1212.5701*, abs/1212.5, 2012. doi: http://doi.acm.org.ezproxy.lib.ucf.edu/10.1145/1830483.1830503.

[30] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *Neural Information Processing Systems*, 2017.

[31] Lee R. Dice. Measures of the Amount of Ecologic Association Between Species. *Ecology*, 26(3):297–302, 1945. doi: 10.2307/1932409.

[32] Shuchao Pang, Zhezhou Yu, and Mehmet A. Orgun. A novel end-to-end classifier using domain transferred deep convolutional neural networks for biomedical images. *Computer Methods and Programs in Biomedicine*, 140:283–293, 2017. doi: 10.1016/J.CMPB.2016.12.019.

[33] Sergi Valverde, Mostafa Salem, Mariano Cabezas, Deborah Pareto, Joan C. Vilanova, Lluís Ramió-Torrentà, Àlex Rovira, Joaquim Salvi, Arnau Oliver, and Xavier Lladó. One-shot domain adaptation in multiple sclerosis lesion segmentation using convolutional neural networks. *NeuroImage: Clinical*, 21:101638, 2019. doi: 10.1016/J.NICL.2018.101638.

[34] Github, NIC-VICOROB/stroke-mri-segmentation, 2019. URL `https://github.com/NIC-VICOROB/stroke-mri-segmentation`.