Computerized Medical Imaging and Graphics 90 (2021) 101908

Contents lists available at ScienceDirect

Computerized Medical Imaging and Graphics

journal homepage: www.elsevier.com/locate/compmedimag

Hemorrhagic stroke lesion segmentation using a 3D U-Net with squeeze-and-excitation blocks

Valeriia Abramova [a], Albert Clèrigues [a], Ana Quiles [c], Deysi Garcia Figueredo [c], Yolanda Silva [b], Salvador Pedraza [c], Arnau Oliver [a,*], Xavier Lladó [a]

[a] Computer Vision and Robotics Group, University of Girona, Catalonia, Spain
[b] Department of Neurology, Hospital Universitari Dr Josep Trueta - Institut d'Investigació Biomèdica de Girona, Girona, Catalonia, Spain
[c] Department of Radiology, Hospital Universitari Dr Josep Trueta - Institut d'Investigació Biomèdica de Girona, Girona, Catalonia, Spain

ARTICLE INFO

Keywords:
Hemorrhagic stroke
Segmentation
Deep learning
Artificial intelligence

ABSTRACT

Hemorrhagic stroke is the condition involving the rupture of a vessel inside the brain and is characterized by high mortality rates. Even if the patient survives, stroke can cause temporary or permanent disability depending on how long blood flow has been interrupted. Therefore, it is crucial to act fast to prevent irreversible damage. In this work, a deep learning-based approach to automatically segment hemorrhagic stroke lesions in CT scans is proposed. Our approach is based on a 3D U-Net architecture which incorporates the recently proposed squeeze-and-excitation blocks. Moreover, a restrictive patch sampling is proposed to alleviate the class imbalance problem and also to deal with the issue of intra-ventricular hemorrhage, which has not been considered as a stroke lesion in our study. Moreover, we also analyzed the effect of patch size, the use of different modalities, data augmentation and the incorporation of different loss functions on the segmentation results. All analyses have been performed using a five fold cross-validation strategy on a clinical dataset composed of 76 cases. Obtained results demonstrate that the introduction of squeeze-and-excitation blocks, together with the restrictive patch sampling and symmetric modality augmentation, significantly improved the obtained results, achieving a mean DSC of 0.86 ± 0.074, showing promising automated segmentation results.

1. Introduction

Stroke is one of the most common causes of death, holding the third position after ischemic heart disease and neonatal disorders (Roth et al., 2018). It is a medical condition in which the brain tissues lose the ability to get oxygen due to a reduced or fully cut blood flow. This rapidly leads to the death of brain cells. There are two types of stroke: ischemic and hemorrhagic. Ischemic stroke is the most common type of stroke (around 87% of all strokes (Mozaffarian et al., 2016)) which is caused by a reduction of the blood supply to the brain tissues; other strokes are hemorrhagic, and they involve the rupture of a vessel inside the brain. In this case, brain cells get damaged due to the pressure of the leaked blood. Even though hemorrhagic stroke is less common, it is characterized by high mortality rates (Kidwell and Wintermark, 2008). Time is the key factor in successful treatment of stroke, since early stroke diagnosis and treatment are related to positive patient outcome (Matsuo et al., 2017). Therefore, fast clinical actions are required in order to give the patient the most appropriate treatment.

The most common imaging modalities for stroke diagnosis are Computed Tomography (CT) and Magnetic Resonance Imaging (MRI), CT imaging being the dominant modality for diagnosing hemorrhagic stroke (Heit et al., 2017). CT is widely available, inexpensive, rapid, and suitable for all patients. In general, when a patient arrives, CT is the first imaging performed. Subsequently, MRI can be used to diagnose other sub-types of hemorrhages.

Hematoma segmentation plays an important part in clinical workflow. It allows tracking the hematoma volume growth quantitatively across several scans, and hence, detecting whether the hemorrhage is still actively bleeding. In addition, the segmentation can help to define a region of interest (ROI) and extract features that could be used for outcome prediction algorithms. In clinical practice the standard approach nowadays is manual delineation of the stroke lesion. However, this approach has disadvantages: it is both time-consuming and operator-dependent, which leads to subjective and non-reproducible results. To address these issues, automated segmentation algorithms have been suggested in recent years (Shahangian and Pourghassem,

* Corresponding author.
E-mail address: aoliver@eia.udg.edu (A. Oliver).

https://doi.org/10.1016/j.compmedimag.2021.101908
Received 8 November 2020; Received in revised form 22 March 2021; Accepted 26 March 2021
Available online 14 April 2021

2015). The initial attempts to segment hemorrhagic stroke lesions in CT images mostly relied on unsupervised approaches such as clustering techniques (Lončarić et al., 1995; Cosić and Lončarić, 1997), morphological operations (Lončarić et al., 1995; Perez et al., 2007), region growing or level sets (Bardera et al., 2009).

In clinical practice, such automated tools can be used as complements in order to support a physician's decisions. For instance, evaluation of hematoma volume is significant and it has a strong independent association with outcome (De Oliveira Manoel et al., 2016), as it allows for improved prognosis when it is combined with physical examination to calculate the ICH (Intracranial Hemorrhage) score (Hakimi and Garg, 2016). Besides hematoma volume, this score is based on other parameters, both patient-related factors and neuro-imaging findings (Hemphill et al., 2001). Nowadays, in clinical practice the ABC/2 method is commonly used to calculate hematoma volume, where A, B and C are dimensions derived from hemorrhage measurements from NCCT (non-contrast CT) images, assuming hematomas to be elliptical shaped (Chinda et al., 2018; Kidwell and Wintermark, 2008). Therefore, this quantification method has some limitations (including processing time, which is very important as stroke is a medical emergency) and introduction of automated tools into clinical workflow could provide a necessary improvement in decision-making speed. Nowadays, commercially available solutions already exist and they help to optimize clinical workflow by accelerating the processing time and by improving diagnostic accuracy (Soun et al., 2020; Wismüller and Stockmaster, 2020).

The recent breakthrough and popularity of deep learning techniques has increased research interest and the number of proposed algorithms, especially Convolutional Neural Networks (CNN), which can be employed to segment stroke lesions. The proposed approaches included 2D-methods, such as the work of Hssayeni et al. (2020), who used a 2D U-Net for segmenting intracranial hemorrhage. Chang et al. (2018) proposed a mask R-CNN algorithm, which used a custom hybrid 3D/2D variant of the feature pyramid network as a backbone to generate a shared set of image features for segmentation of different types of hemorrhagic stroke. 3D approaches were also proposed. For instance, Singh et al. (2019) presented a 3D CNN to segment several hematoma types, employing a novel thresholding method as a pre-processing step, while Hu et al. (2020) used an encoder-decoder CNN for segmenting intracranial hemorrhages. A patch-based 3D U-shaped network was proposed by Patel et al. (2019), although cases with sub-arachnoid and intra-ventricular hemorrhages were excluded from their dataset. The class imbalance issue was addressed by utilizing weighted maps for each patch. U-shaped architectures were also applied for hemorrhagic stroke (Wang et al., 2018). Kwon et al. (2019) used a Siamese U-Net with a healthy template to segment different sub-types of hemorrhagic stroke. Yao et al. (2020) proposed the use of multi-view CNN with a mixed loss function. This architecture shared some similarities with a U-Net, while the mixed loss function was designed to make the system robust to CT scans acquired in different medical centers, using different protocols. Other architectures such as a cascade CNN were built to detect and delineate several hemorrhage sub-types (Cho et al., 2019), including intra-ventricular hemorrhage. Kuang et al. (2020) presented a network named ψ-Net, where two attention blocks were used to suppress irrelevant information, and capture the spatial contextual information to refine the contours of the lesion. The method was evaluated on 2D-slices with different hematoma types. Finally, Kuo et al. (2019) proposed a patch-based FCN with Dilated ResNet 38 backbone, which performed both classification and segmentation tasks.

In the present work, a deep learning approach for segmentation of hemorrhagic stroke lesions in non-contrast CT is proposed. In contrast to previous works, we introduced squeeze-and-excitation blocks to a U-Net architecture, similarly to the work of Woo et al. (2019) for ischemic stroke segmentation, but, in our case, using a 3D U-Net implementation. A balanced sampling technique is also proposed for patch extraction, restricting the sampling spatially and quantitatively to address the problem of intra-ventricular hemorrhage (i.e. blood flowing into the ventricles). To our knowledge, this is the first approach that deals with intra-ventricular hemorrhage as a background class, focusing only on the stroke lesion segmentation. Previous works utilized datasets either without this sub-type of stroke or having it segmented as another stroke class. In addition to standard data augmentation techniques, symmetric modality augmentation is also introduced to benefit from the brain hemispheres symmetry property to find more robust image features (Clèrigues et al., 2019). Our segmentation proposal is evaluated with a cross-validation strategy over a clinical dataset of 76 cases, analyzing in detail the impact of the different contributions introduced in our approach.

2. Materials and methods

2.1. Dataset

The dataset used in this project consisted of 76 non-contrast head CT images acquired in the Hospital Dr. Josep Trueta, Girona, Spain. Some of the cases included also CT angiography and CT perfusion images. All the examinations were performed on a 128-slice CT scanner (Ingenuity; Philips Healthcare), whose image characteristics are summarized in Table 1. For non-contrast CT (CT NC) the slice thickness was 3 mm and the gap was 1.5 mm, whereas for CT angiography (CT Angio) the slice thickness was 0.9 mm with a gap of 0.45 mm. CT perfusion images consisted of 4 slices of 10 mm (Puig et al., 2017), but were not used in our work. The upper row of Fig. 1 shows an example of a non-contrast CT and CT angiography image used in our work.

The gold reference standard consisted of the masks of the hemorrhagic stroke lesions manually delineated by expert radiologists on the non-contrast CTs.

2.1.1. Intraventricular hemorrhage

Intra-ventricular hemorrhage (IVH) or intra-ventricular bleeding is an extension of hemorrhage, which occurs within brain parenchyma, inside brain ventricles, where the cerebro-spinal fluid is produced. One of its sources can be the hemorrhagic stroke lesion adjacent to the ventricles. Since the source of signal in both pathologies is the same (blood leakage), the intensities of both regions on non-contrast CTs are also similar (Fig. 2). Such pathology is a bad prognosis sign, as expected mortality from it is between 50% and 80% (Hinson et al., 2010). Usually, IVH can be clearly confirmed from CT imaging from the presence of blood inside the ventricles. Therefore, the problem of segmenting intra-parenchymal hemorrhage (IPH) from IVH is not simple.

Almost 15% of the cases in the Trueta dataset had IVH together with the stroke lesion. As the primary objective was the stroke lesion segmentation, located only within brain tissues, the intra-ventricular hemorrhage was not delineated as a groundtruth. Thus, the developed method needs to be able to differentiate between IPH and IVH. As we will show later, the IVH is considered as part of the background class.

2.1.2. Data preparation

The initial preparation of non-contrast CT head scans requires removal of the coil and skull stripping (see Fig. 1) as those regions can confuse the algorithm and lead to undesired results. Both coil removal and skull stripping were done using morphological operations, like extracting the biggest connected component for coil removal. The skull

Table 1
Image characteristics of the dataset modalities used in our work: non-contrast CT (CT NC), CT angiography (CT Angio).

Image modality	Matrix	Slice thickness, mm	Gap, mm
CT NC	512 × 512	3	1.5
CT Angio	512 × 512	0.9	0.45

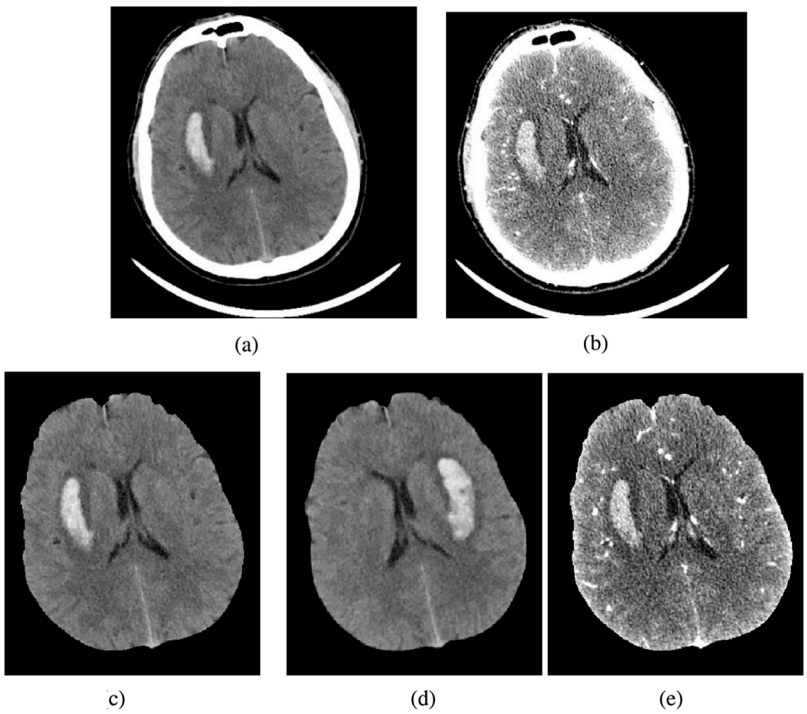

Fig. 1. (a-b) Original image modalities and (c)–(e) pre-processed images used as input to the U-Net. (a) Non-contrast CT and (b) CT angiography, (c) pre-processed non-contrast CT, (d) symmetric version of (c), and (e) pre-processed CT angiography.

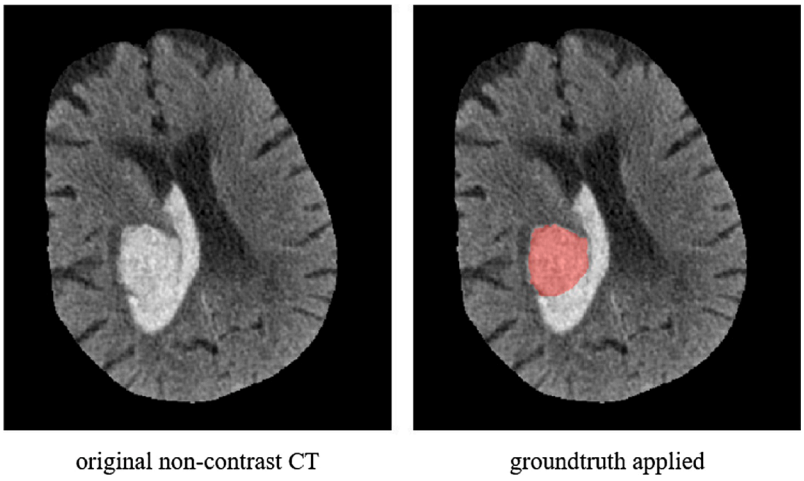

Fig. 2. An example of intra-ventricular hemorrhage. Both IPH and IVH have similar intensities, even though the groundtruth provided only has IPH. Moreover, IVH deforms brain ventricles.

removal process is similar, it utilizes morphological operations to remove the borders of the skull and the final brain extraction is also based on extracting the biggest connected component.

As stroke can appear in only one of the hemispheres of the brain, it could be useful to utilize features based on the mid-sagittal symmetry of those hemispheres, as proposed in Clèrigues et al. (2019). Therefore, a symmetric image of the brain was obtained from each image as follows. We flipped the brain CT and subsequently we registered this flipped image to the initial one using the FLIRT algorithm (from the FSL toolbox (Jenkinson and Smith, 2001; Jenkinson et al., 2002)). The overall pre-processing time including creation of symmetric modality took on average 4 min per image. Fig. 1 also shows one example of a flipped image.

Additional pre-processing is necessary to use the angiography images. We cropped some images for cases that contained the upper body, and also registered the images to the non-contrast CT space so as to have the same voxel spacing. The final output of this processing is shown in the bottom row of Fig. 1.

2.2. Proposed method

Our proposed approach for segmenting the hemorrhagic stroke lesion is a 3D patch-based deep learning method based on a U-Net architecture. As the lesion mostly occupies only a small volume inside the brain, class imbalance is a problem that has to be taken into account when training the network to avoid over-fitting to the negative class. Moreover, the dataset used assigned the IVH to background and therefore this has to be tackled within the developed algorithm. To address these issues we propose: (a) a balanced sampling technique to ensure equal distribution of both classes in the training set, and (b) restrictions

in regions to extract patches to distinguish between intra-parenchymal and intra-ventricular hemorrhages. At the training stage, to tackle these problems, regularization techniques are also applied, such as: (a) dropout, (b) data augmentation, and (c) early stopping. During testing, high overlap between extracted patches was used to improve the segmentation results.

2.2.1. Patch sampling

The proposed approach is a patch based CNN architecture, which alleviate from computationally heavy load of large input images and also offers reduced training time (Long et al., 2015). For patch based methods, class imbalance can be an issue so it is important to control the process of patch extraction. In our work we used a balanced sampling strategy, ensuring that an equal number of patches representing both classes (lesion vs background) were extracted from each image. To avoid extracting a lot of patches from background and take more advantage of the dataset, the area to extract non-lesion patches was limited within the brain mask; in this way negative patches were uniformly extracted only from the region inside the brain. However, similarly to the work of Clèrigues et al. (2019), we utilized additional ways of controlling the patch extraction for specific aims.

With the aim of improving the segmentation results on lesion borders, we extracted additional negative samples from the region around the hematoma border. Thus, we make the network learn more from these boundaries. This approach was inspired by Kushibar et al. (2018), who successfully used it for the segmentation of brain sub-cortical structures.

As mentioned before, some images of the dataset contained intra-ventricular hemorrhage, which is not considered part of the intra-parenchymal lesion, but has a similar appearance. The first step to resolve this issue is taken during data preparation, when sampling the training patches. Voxels of IVH belong to the negative class, whereas the voxels of IPH belong to the positive (lesion) class. To make the network learn this dependence, we ensure that patches from the IVH area are represented in the training set. Since this abnormality occurs inside the brain ventricles, we use spatial information related to brain anatomy and we extract additional patches from the area of CSF ventricles. In the experimental section we will test two approaches for solving this issue. The first one comes from the observation that the brain ventricles are located in the medial area of the brain. Hence, we define a region of interest around the center of the brain volume. The coordinates of the brain center were calculated for each image separately. For each dimension we defined the first and last slice where the brain appears and took the coordinates of the middle slice. On the other hand, the second approach comes from the fact that intra-ventricular hemorrhage are hyper-intense voxels inside the brain (that do not belong to IPH). To define it automatically, we thresholded the whole image using an empirically found threshold that represented the intensity of the blood signal and then we excluded those voxels belonging to the lesion by using the available groundtruth. Moreover, some other hyper-intense structures such as brain borders and midlines were excluded by utilizing the fact that IVH has a bigger volume. Therefore, with the ROI defined in one of these two ways, we restrict the training patch sampling by assigning the fraction of negative patches forcibly extracted from this region. As the number of scans with IVH is not that huge, we may extract more training patches from this region by extracting the background patches only from the defined ROI.

Hence, in the overall patch extraction pipeline, firstly, the target number of patches was set for each patient. 50% of patches were extracted uniformly from the whole volume of the brain, including a predefined fraction of them that were forcibly extracted from the area around the lesion boundary and the area related to the brain ventricles. Here, uniform sampling was done in order to make sure that all the parts of the brain were equally represented. In addition, 50% of patches were extracted from the lesion voxels.

Furthermore, data augmentation techniques can be applied to the extracted patches, increasing the size of the patch dataset proportionally to the number of patches specified at the beginning. Reasonable transformations are applied in terms of brain anatomy: horizontal and vertical flip, and patch rotations of 90°, 180°, and 270°. Therefore, in the training set, one patch was used six times, each time transformed with a different augmentation type.

The patches were extracted from non-contrast CT scans but as we will see in Section 3, some experiments were complemented with their symmetric versions or angiography CT as additional input channels to the network.

2.2.2. U-net with squeeze-and-excitation blocks

As shown in Fig. 3, our work is based on a 3D U-Net architecture which incorporates squeeze-and-excitation blocks. We chose the U-Net considering its prevalence in medical imaging segmentation and the fact that it has been proven to work reliably with small datasets, while 3D patches provide more information than 2D. Similarly to what was done in the work of Woo et al. (2019) for the ischemic stroke segmentation problem, we propose to incorporate squeeze-and-excitation blocks into the 3D U-Net architecture, as they have shown improved performance for segmentation tasks. We use single $3 \times 3 \times 3$ convolutions in both contracting and expansive paths and rectified linear unit (ReLU) as activation function. Down-sampling in each resolution step is performed by a $2 \times 2 \times 2$ max pooling operation with a stride of 2, followed by dropout. In the expansive path, the up-sampling is performed with up-convolution of $2 \times 2 \times 2$ with a stride of 2, followed by dropout. We introduce squeeze-and-excitation operations after encoder and decoder, as shown in Fig. 3. The goal of these blocks is to improve the quality of representations produced by a network by explicitly modeling the inter-dependencies between the channels. Structurally, this computational unit consist of: (a) a *squeeze* operator, which produces a descriptor for each of the channels of the image; and (b) an *excitation* operator, which aims to fully capture channel-wise dependencies. All this is done to enhance the informative features and suppress the weak ones. The effect of including these blocks will be evaluated in the experimental results section. The total number of parameters of the proposed network is 2,263,042. The overall configuration of network training and testing is described in the following section.

2.2.3. Training and testing pipelines

In the training stage, we composed the training and validation sets from the provided scans to train the weights of the network. As we will see in the experimental section, the effect of different numbers of patches and different patch sizes is analyzed. At the end, from each image in the training set $3,000$ patches of size $32 \times 32 \times 16$ were extracted following the patch sampling technique introduced previously, as presented in Fig. 4. Notice that the positive patches, shown with a green contour, are centered on lesion voxels, while negative patches, with red contours, are centered on the voxels not related to IPH.

To take advantage of symmetry features, the symmetric image was added as another channel to each training image. Data augmentation was performed to increase the number of patches per image to $18,000$, and its effect will be also analyzed in Section 3.

Two different loss functions were also studied: focal loss and combined Dice and Cross-entropy losses. Focal loss (Lin et al., 2017) gives less weight to easily classified examples and more weight to hard to classify examples, therefore it can be useful for the tasks involving class imbalance. On the other hand, a combination of Dice loss and Cross-entropy loss has also been popular for ischemic stroke segmentation tasks (Clèrigues et al., 2019). Adadelta was used as an optimizer since it did not require manual tuning of the learning rate (Zeiler, 2012). To prevent over-fitting, an early stopping technique with a patience tolerance of 15 was utilized when approaching the minimal loss on validation set. In addition, either the maximum number of epochs to train was set to 100 or else the training was performed until it met the early stopping condition. In practice, the number of epochs to train was

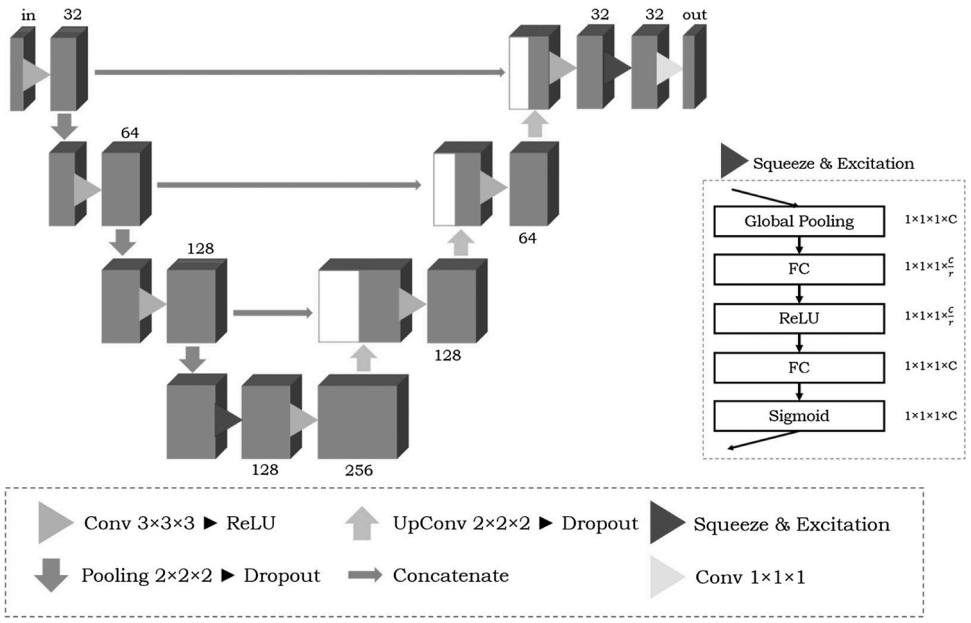

Fig. 3. The 3D architecture used in the proposed approach. The network is inspired by a 3D U-Net with incorporation of squeeze-and-excitation blocks. Blue blocks represent feature maps with a number of channels stated above or below them. Red triangles represent the location, where we incorporate squeeze-and-excitation operations. In the detailed description of these operations, C and r represent the number of input channels and reduction ratio, respectively.

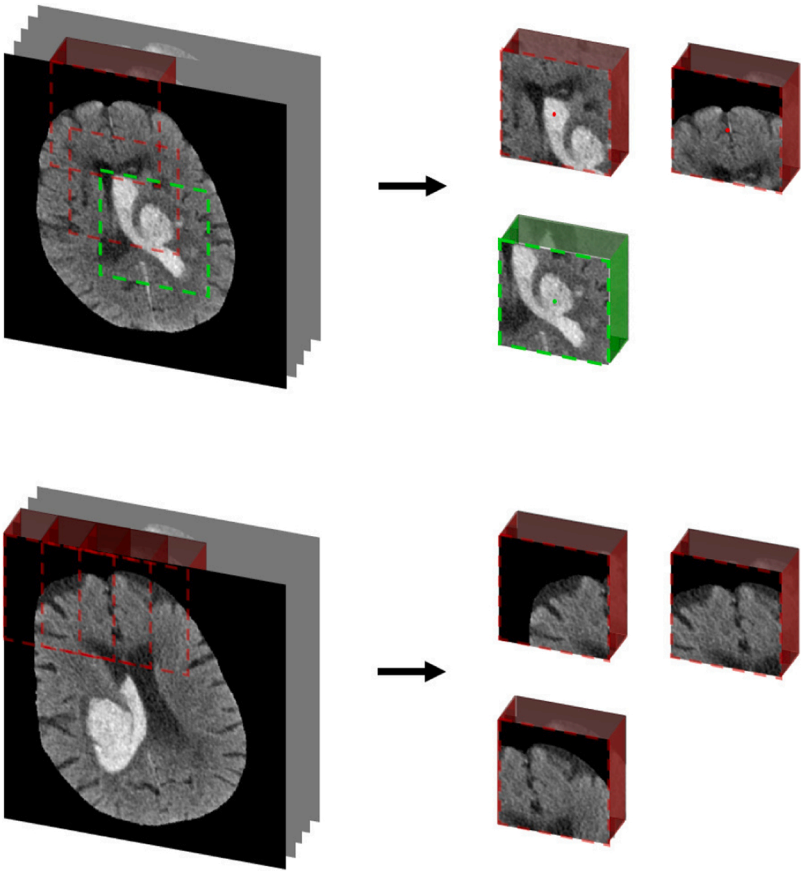

Fig. 4. Implemented patch sampling for training and testing stages. During training (upper figure), negative patches (red contour) are extracted uniformly from the whole volume of the brain, and some fraction of them is extracted from the area around a lesion boundary and the area related to the brain ventricles. Positive patches (green contour) are centered on lesion voxels. During testing (bottom figure), the patches are extracted uniformly from the whole volume, preserving some degree of overlap.

between 20 and 50.

For the testing stage, given an image to segment, patches of the same size as in the training step were extracted uniformly from the whole image volume, as shown in Fig. 4. The extracted patches preserve a predefined degree of overlap between them, more than 50%, which was done to improve segmentation results (Bernal et al., 2019). Every patch was passed through the network, resulting in a predicted probability for each voxel. The output binary segmentation was produced by assigning the class label according to the maximum probability for each voxel.

2.2.4. Implementation details

The proposed approach was implemented in Python using the Pytorch machine learning framework (Paszke et al., 2017). All experiments were run on Ubuntu with 256GB RAM, and the network training was done on TITAN V GPU with 12 GB of memory.

2.3. Evaluation

All the experiments were performed with a 5-fold cross-validation across all 76 cases of the provided dataset, made up of 61 images in the training set and 15 images in the testing set.

The Dice similarity coefficient (DSC) was used as an evaluation metric, as it is widely used to assess segmentation tasks as a measure of overlap between output segmentation and groundtruth:

$$\mathrm{DSC} = \frac{2\mathrm{TP}}{2\mathrm{TP} + \mathrm{FP} + \mathrm{FN}}$$

where TP, FP and FN refer to true positive, false positive and false negative voxels, respectively.

To evaluate the statistical significance of differences between the obtained results, we used the dependent t-test for paired samples.

3. Experimental results

Experiments were performed to show the improvement of the pipeline through the development process. We analyzed the influence on the segmentation results for the different steps involved in the pipeline: 1) incorporation of squeeze-and-excitation blocks to the baseline U-Net architecture, 2) loss functions, 3) restrictive patch extraction to improve segmentation and to deal with the intra-ventricular hemorrhage, 4) use of different modalities, and 5) impact of data augmentation.

3.1. Squeeze-and-excitation blocks

The first experiment performed compared the performance of the proposed U-Net with the standard architecture. A balanced patch sampling technique was used, extracting 3000 patches of size $(32, 32, 16)$ per patient, which empirically provided the best trade-off between performance and computational cost. The negative patches were extracted from the brain area only. The results are summarized in Table 2, detailed for the cases with and without intra-ventricular hemorrhages.

From the table, we can see that incorporating the squeeze-and-excitation blocks into the standard U-Net architecture significantly improved the overall segmentation results of the dataset, increasing the average DSC and reducing the standard deviation ($p < 0.01$), although it does not help to break through the maximum segmentation DSC obtained, changing only from 0.967 to 0.968. A qualitative comparison between using both architectures is shown in Fig. 5.

3.2. Loss functions

The second experiment relates to the performance of the loss functions analyzed, focal loss and combined Dice and cross-entropy loss. The initial conditions for the experiment were using 3000 patches of size $(32, 32, 16)$ per image, and using brain mask and ROI around hematoma as restrictive conditions.

The resulting segmentations of the model trained with the focal loss showed an average DSC of 0.796 ± 0.158, while using the combined cross-entropy and Dice loss, we obtained an average DSC of 0.841 ± 0.108, improving the results significantly ($p < 0.01$).

Table 2
DSC obtained in the dataset in cross-validation experiment with and without incorporation of squeeze-and-excitation blocks (SE) into the 3D U-Net. The best results are highlighted in bold.

	DSC	DSC of IVH samples	DSC of samples with no IVH
3D U-Net	0.765 ± 0.217	0.640 ± 0.217	0.787 ± 0.212
3D SE U-Net	**0.828–± 0.127**	**0.683–± 0.118**	**0.852–± 0.112**

3.3. Patch size and restrictive patch extraction

The third experiment performed was the study of the effect of the patch size and the sampling strategies. Considering the architecture used (with squeeze-and-excitation blocks incorporated) and the computational load, three patch sizes were tested: $(24,24,8)$, $(32,32,16)$, $(48,48,24)$. The obtained results are visually presented in the boxplots of Fig. 6. On the one hand, increasing patch size to $(32, 32, 16)$ significantly improved mean DSC ($p < 0.001$) from 0.759 ± 0.183 to 0.842 ± 0.115, but, on the other hand, when the patch size was enlarged to $(48,48,24)$, the average DSC was significantly decreased ($p < 0.01$) to 0.805 ± 0.156. Moreover, we observed that none of the patch sizes helped to surpass the maximum DSC obtained (being 0.958, 0.968 and 0.963 for each patch size, respectively). From Fig. 6 we noticed that the results obtained with patches of the size $(32, 32, 16)$ were within the smaller range, which implies a lower dispersion. Therefore, for the rest of experiments we used the medium patch size of $(32, 32, 16)$, that provided consistent performances.

We also analyzed the effect of the different restricted sampling strategies, obtaining the results summarized in Table 3. Constraining the area to extract negative patches within the brain mask significantly improved segmentation results ($p < 0.001$). To improve and refine the lesion contours, we also defined a region of interest around the hematoma and a fraction of patches, which we defined empirically to be 30%, was extracted from this region. This ROI was represented as a cubic volume around the hematoma borders. Defining this region allowed the network to focus more on the area near the lesion borders and helped to significantly improve the obtained segmentation ($p < 0.01$). In addition, with this experiment, the minimum DSC obtained per case was notably improved, going from a DSC of 0.183 to 0.530.

To tackle the problem of differentiating intra-parenchymal and intra-ventricular hemorrhage, we forced the network to learn more from IVH areas. Two different approaches were proposed: to define a cubic region of interest (ROI) around the central area of the brain or to define the ROI by just considering the IVH area (i.e. hyperintense voxels). The results of this experiment are shown in the two bottom lines of Table 3. Using the first approach, results increase by a small margin, although the standard deviation was also increased, meaning that more variability was added. The improvement was not significant ($p > 0.05$). Surprisingly, the improvement was larger in cases without IVH. In contrast, with the second approach, the results significantly decreased ($p < 0.01$).

3.4. Usage of additional input channels

The fourth experiment deals with the use of additional channels as input to the net. The first channel added is the symmetric non-contrast CT image, with the aim of exploiting brain symmetry. The experiment was performed, maintaining the same training parameters and patch sampling strategies defined above. The obtained results are presented in Table 4. Notice that adding the symmetric image with the fixed ROI around the lesion improved the mean DSC from 0.841 ± 0.108 to 0.849 ± 0.099, being statistically significant ($p < 0.01$). The improvement was also observed when segmenting the images with IVH, with a DSC increase of 2.8%, while also reducing the standard deviation. Qualitatively, this can be seen in Fig. 7, where yellow arrows indicate the areas that were correctly segmented as background after introducing the symmetric input. When extracting patches considering the cubic ROI around the brain center and using a symmetric image as additional input channel the segmentation results improved, although the improvements

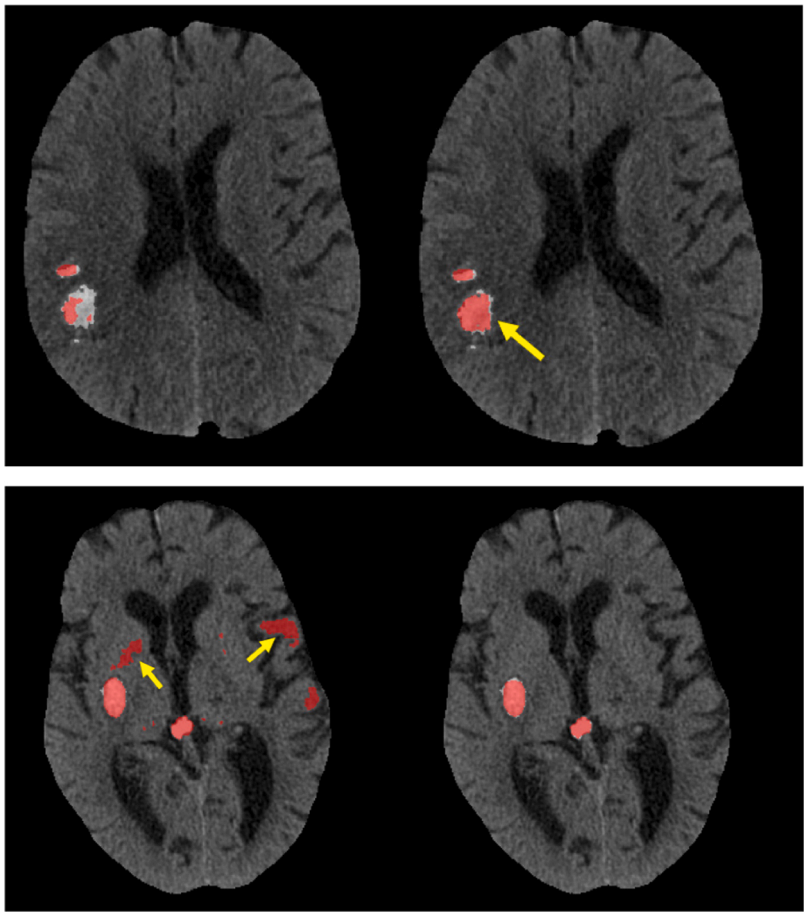

Fig. 5. Two examples of the benefits of introducing squeeze-and-excitation blocks (right images) into the baseline 3D U-Net (left images). The yellow arrows show the areas of improvement. The upper image shows a better segmentation of the lesion borders, while the bottom image shows that introducing squeeze-and-excitation blocks helps in removing false positives lesions.

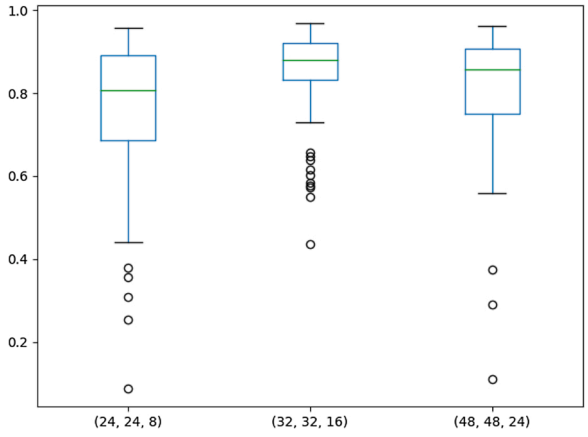

Fig. 6. DSC values obtained within three experiments with different patch sizes: $(24, 24, 8)$, $(32, 32, 16)$, $(48, 48, 24)$.

were not significant ($p > 0.05$). However, if symmetric modality augmentation was used and patches were forcibly extracted from IVH volumes, the obtained DSC results, especially for the images with IVH (see Table 4), significantly improved ($p < 0.01$) with an increase of 4.9%.

We also analyzed the inclusion of angiography CT as an additional input channel together with the original non-contrast CT. For a fair comparison, the cases without angiography were excluded (10 cases in total). However, we observed that when introducing angiography, the DSC significantly decreased ($p < 0.001$) from an average of 0.834 ± 0.120 to 0.795 ± 0.156.

3.5. Data augmentation

Finally, we studied the influence of data augmentation by applying the transformations presented in Section 2.2.1 to all the patches extracted within the experiment (i.e. for each patch, 5 additional patches were obtained). Results are shown in Table 5. Firstly, the number of extracted patches was reduced to 500 and applying data augmentation was augmented up to 3000. Compared to the baseline approach (first line of the Table), the segmentation results significantly decreased ($p < 0.05$). However, when increasing the number of extracted patches, the results applying data augmentation slightly improved, only decreasing when the initial number of patches was 1500. For the other cases, the changes were statistically significant ($p < 0.01, p < 0.01$ and $p < 0.005$, respectively). However, the difference between the results obtained with 3000 patches and those with 3000 patches augmented to $18,000$ was not statistically significant ($p > 0.05$).

3.6. Final configuration

Taking into account all the previous experiments, the final configuration of the proposed method was chosen. The network architecture included squeeze-and-excitation blocks and was trained with 3000 patches of size $(32, 32, 16)$ using as a loss function the combination of Dice and cross-entropy losses. Even though results with data augmen-

Table 3
The resulting DSC for the whole dataset and its parts with and without hemorrhage using different patch restriction steps. The values in bold specify the approach with the highest metric value. The significance of adding each patch restriction step is shown with * and **.

	DSC	DSC of samples with IVH	DSC of samples without IVH	max DSC	min DSC
No restrictions	0.599±0.284	0.546±0.239	0.608±0.291	0.946	0.028
Brain mask	0.807±0.159**	0.664±0.188	0.831±0.141	0.959	0.183
Brain mask + ROI around hematoma	0.841±0.108*	0.692±0.128	0.866±0.081	0.964	0.530
Brain mask + ROI around hematoma + ROI around brain center	0.842±0.115	0.699±0.126	0.867±0.094	**0.968**	0.435
Brain mask + ROI around hematoma + hyperintense ROI	0.823±0.136*	0.687±0.124	0.846±0.125	0.963	0.306

* Significant at $p < 0.01$.
** Significant at $p < 0.001$.

Table 4
The evaluation metrics of the experiment with and without symmetric modality as an additional input channel. Different patch restriction cases are compared. The values in bold specify the approach with the highest metric value. The significance of adding the symmetric modality as additional channel for each patch sampling strategy is shown with *.

	Input modalities	DSC	DSC of samples with IVH	DSC of samples without IVH	Max DSC	Min DSC
ROI around hematoma	Original image	0.841±0.108	0.692±0.128	0.866±0.081	0.964	0.530
	With symmetric modality	0.849±0.099*	0.720±0.116	0.871±0.078	0.964	0.530
ROI around hematoma + brain center	original image	0.842±0.115	0.699±0.126	0.867±0.094	**0.968**	0.435
	With symmetric modality	0.856±0.088	0.728±0.112	0.878±0.061	0.956	0.553
ROI around hematoma + IVH area	original image	0.823±0.136	0.687±0.124	0.846±0.125	0.963	0.306
	With symmetric modality	0.857±0.085*	0.728±0.107	**0.879±0.057**	0.964	0.581
ROI around hematoma OR IVH area	with symmetric modality	**0.862±0.074**	**0.777±0.101**	0.876±0.059	0.957	**0.632**

* Significant at $p < 0.01$.

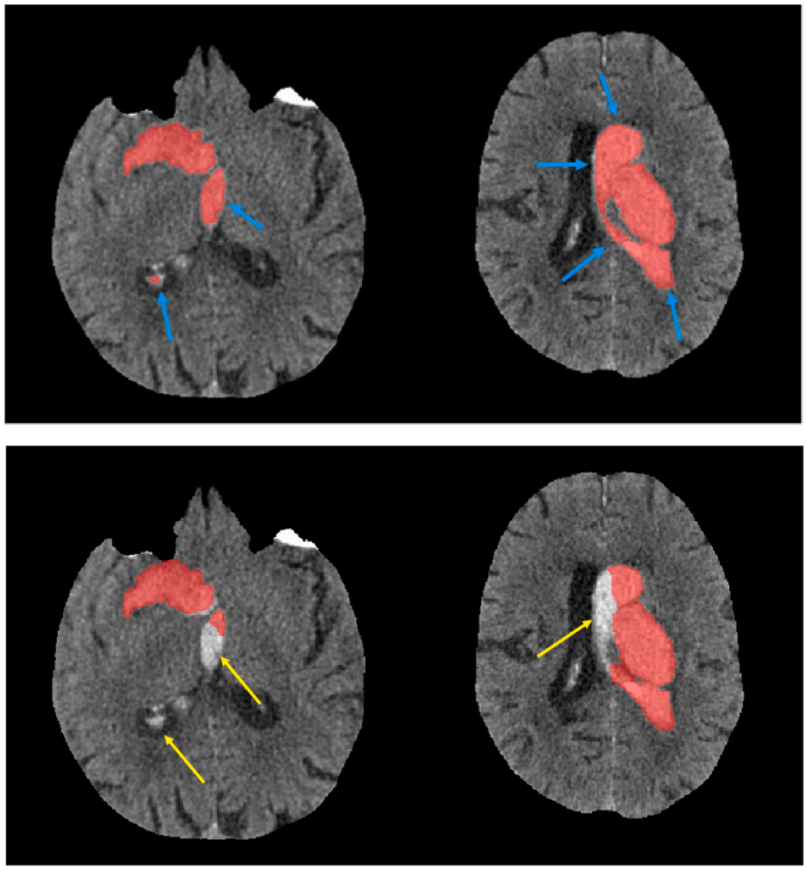

Fig. 7. Results of hematoma segmentation without (top) and with (bottom) the symmetric image as an additional input channel. Blue arrows indicate the IVH areas, that are wrongly segmented without using the symmetric image. Yellow arrows show the changes in the results achieved by incorporating the symmetry information.

Table 5
The evaluation metrics of experiments checking the influence of data augmentation and its size on the segmentation results.

	Number of patches	DSC	Std
Augmentation	3000	0.862	0.074
	500 → 3000	0.845	0.092
	1000 → 6000	0.866	0.078
	1500 → 9000	0.846	0.092
	3000 → 18000	0.868	0.076

tation improved the mean DSC, considering the computational cost *versus* the amount of improvement, data augmentation was not considered in the final approach. The original image together with its symmetric modality were used as input, and patches were extracted from them restrictively, limiting the area within the brain mask to extract patches, to the ROI around hematoma, and also the IVH volume. With this final design, we could achieve a mean DSC of 0.862 with an approximated segmentation time of 17 seconds per patient. For this final network configuration the average training time per one epoch was 5 min 33 s. The qualitative example of one of the segmentation results is shown in Fig. 8.

The quantitative comparison with the current state-of-the-art approaches is difficult for several reasons. Firstly, all state-of-the-art methods used different datasets (most of them private), which may include different levels of disease severity. Secondly, our goal was to segment intra-parenchymal hemorrhage only, however, some images in our dataset included cases with intra-ventricular hemorrhage, and our aim was to avoid segmentation of these areas. In contrast, in other state-of-the-art approaches, IVH cases were either not presented at all, or were excluded, or segmented as another class. This means that these works could train specifically IPH regions only, providing higher overall DSC values. Some of the proposed methods used datasets with more scans than our dataset had, which could also influence the resulting segmentations. Table 6 shows the summary of all reviewed state-of-the-art approaches. For instance, the work of Chang et al. (2018), reported an average DSC of 0.931 with a dataset of 10159 CT scans where 8.9% of scans had IVH, while the approaches of Singh et al. (2019) and Kuang et al. (2020) provided mean DSC values of 0.932 and 0.864, with datasets composed of 399 and 150 scans, respectively. In the work of Patel et al. (2019) IVH cases were excluded, and they were able to achieve median DSCs of 0.91 and 0.90 on datasets of 51 and 100 scans. The approach of Yao et al. (2020) reported an average DSC of 0.697 with a dataset of 120 CT scans from different centers. One of the recent approaches by Sharrock et al. (2020) compared several configurations of

Table 6
Comparison between the proposed approach and existing deep learning methods.

Author	Method	Dataset	Groundtruth	DSC for IPH
Proposed method	3D SE U-Net	76 3D CTs	Annotated IPH masks, IVH is considered as background	0.862
Wang et al. (2018)	3D CNN	243 3D CTs	.csv files with hematoma annotation	qualitative results
Chang et al. (2018)	Mask R-CNN	10159 CT scans for training; 8.9% with hemorrhage	3D masks with 3 stroke subtypes delineated	0.931
Singh et al. (2019)	3D CNN + thresholding	399 3D CTs	3D masks with 4 stroke subtypes delineated	0.932
Hssayeni et al. (2020)	2D U-Net	82 3D CTs; 36 with ICH	3D masks with 5 stroke subtypes delineated	0.280
Kuang et al. (2020)	ψ-Net	150 CT scans	Annotated ICH masks	0,950 0,894
Yao et al. (2020)	Multi-view CNN	120 CT scans	Annotated ICH masks	0.697
Sharrock et al. (2020)	2D and 3D V-Net	112 CT scans	3D masks with 3 stroke subtypes delineated	0.911
Arab et al. (2020)	CNN-DS	55 CT scans	Annotated ICH masks	0.840

V-Net, reporting the best DSC of 0.911. Their training dataset consisted of 112 CT scans with 3 stroke sub-types delineated, including IVH. Finally, the approach of Arab et al. (2020) proposed a CNN with deep supervision, reporting an average DSC of 0.840 with a dataset of 55 scans. Note that most of the analyzed state-of-the-art approaches had more data and, additionally, all stroke sub-types were segmented as a lesion class. Nevertheless, the overall results obtained for IPH with our approach (DSC of 0.879) should be considered as very satisfactory, especially taking into account the challenges presented in our dataset, both the consideration of not including the intra-ventricular hemorrhage as a lesion, and the number of such cases in the whole dataset, being almost 15%.

4. Discussion

This study presented a deep learning approach for hemorrhagic

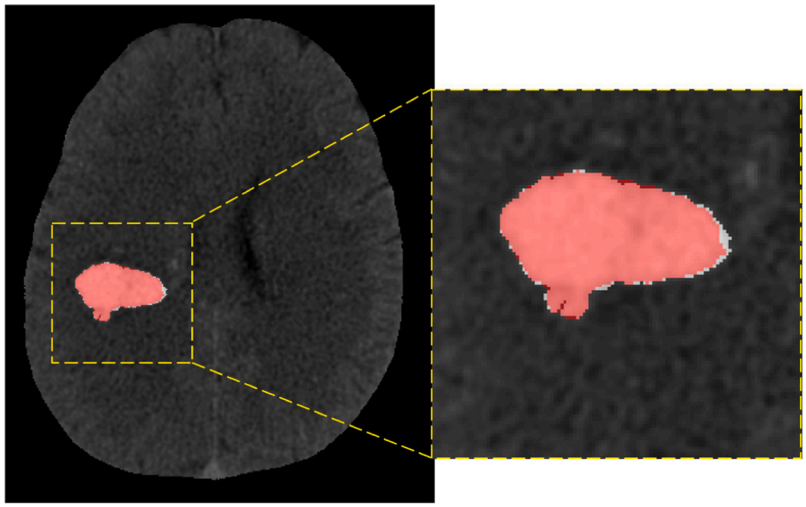

Fig. 8. Example of a good segmentation result. The red overlay represents the resulting segmentation, while the white one represents groundtruth.

stroke lesions segmentation, which appear as hyper-intense regions on CT images. As already shown in previous studies, promising performance could be achieved. Using a baseline patch-based 3D U-Net architecture with the only restriction of using the brain mask for sampling the negative class, we already achieved a mean DSC of 0.765 ± 0.217, as shown in Table 2. Notice that this result is better than the one reported by Yao et al. (2020), who achieved an average DSC of 0.697. Their approach was also inspired by a U-Net, but their dataset was more heterogeneous, with images acquired in different medical centers, and this may be the reason of such a lower performance. In addition, the incorporation of squeeze-and-excitation blocks helped to significantly improve the average segmentation DSC by 6.3%, as well as to reduce variability by reducing the standard deviation from 0.217 to 0.127 (Table 2). Looking at the obtained results, the addition of squeeze-and-excitation operations helped to better segment small and irregular lesions, as can be seen in Fig. 5. This is in line with the results obtained by Woo et al. (2019) when segmenting ischemic lesions. Regarding loss functions, the combination of Dice and cross-entropy losses significantly outperformed the focal loss, increasing the average DSC by 4.5% while reducing the standard deviation by 5%.

The experiments carried out with different patch sizes showed an unexpected behavior. Enlargement of patch size should have helped the network to capture more contextual information, however, our experiments showed that increasing the patch size up to $(48, 48, 24)$ worsened the mean DSC and introduced more variability to the samples, as shown in Fig. 6. Moreover, the large patch size was not able to help to capture dependencies between IPH and IVH. Individually, the most significant decrease of DSC happened with small lesions, which shows that for this particular dataset the patch size of $(48, 48, 24)$ was too big to capture the small hematoma volumes. A similar behavior was observed in the work of Bernal et al. (2019) for tissue segmentation using MRI. In our case, increasing patch size from $(24, 24, 8)$ to $(32, 32, 16)$ led to better segmentations, as expected initially. Therefore, the patch size of $(32, 32, 16)$ was chosen as the optimal one.

A restrictive patch sampling strategy was used. The initial constraint of extracting patches only from the brain volume was performed in order to generally improve the segmentation. The obtained results showed that sampling patch centers only from brain voxels greatly increased the DSC over all cases of the dataset. The significant increase of the DSC in the case of fixing the region of interest around the lesion was also expected, since it was also shown in the work presented by Kushibar et al. (2018). In most cases introducing this condition helped to refine the lesion contours, as can be observed in the qualitative example provided in Fig. 8. However, for some images improper outputs were produced. See for instance the yellow arrows of the qualitative examples shown in Fig. 9. Notice that less patches were extracted from the overall brain volume, therefore some voxels were mistaken for lesions. In contrast, the attempts to solve the issue of intra-ventricular hemorrhage by establishing another region of interest around CSF ventricles did not significantly change the hemorrhage lesion segmentation, introducing more variability to the data (Table 3). This was possibly due to the fact that the network received more patches with IVH and from the ventricles, and the overall distribution of patches was not enough to successfully distinguish between IPH and IVH.

Symmetric modality augmentation showed an interesting behavior. Even without extra guidance to ventricle areas, it could improve segmentation results of both groups of scans, with and without IVH, exploiting the fact that hemorrhage occurs in only one of the brain hemispheres. Even though fixing the ROIs related to brain ventricles did not help to achieve better results, they were successfully used together with symmetric modality augmentation. As can be observed in Table 4, fixing these ROIs in the ventricles improved the results compared to fixing only the ROIs around the hematoma. Notice that these constraints were exploiting the information about ventricle shapes and their deformation in the presence of intra-ventricular hemorrhage. Qualitatively, Fig. 7 shows that incorporating the symmetric modality as an additional input channel could reduce segmentation of voxels located to the other hemisphere of the brain damaged by stroke. Moreover, segmentation of voxels of the same hemisphere, but related to intra-ventricular hemorrhage was also reduced. In contrast, we noticed a decrease in the DSC results after introducing CT angiography as additional input channel. This could be due to the fact that these images were more noisy than non-contrast CT scans.

Regarding the artificial expansion of the dataset using data augmentation techniques, we observed that the results did not provide a breakthrough in the performance of the algorithm. Moreover, augmenting the patches from 3000 to $18,000$ did not introduce much improvement compared to the results obtained using 3000 original patches without any augmentation (see Table 5). The overall difference in the obtained results was not significant ($p>0.05$). Considering the mean DSC and standard deviation values from the different experiments, one can see that the general behavior of the DSC distribution was similar.

Finally, as can be seen from Table 4, the minimum DSC achieved with the final configuration of the proposed approach is 0.632. If we analyze DSCs of all the cases we studied, only 4 have DSC less than 0.7, and 3 out of these 4 have IVH presented, which was expected, as IVH is the main challenge of our dataset. However, from Table 4 we notice that our final approach could maximally increase the minimal DSC if compared with previous experiments. 22% of the dataset is composed of the images with irregular lesions, and for these cases only 3 out of 17 images have DSC < 0.8. Moreover, 14 CT scans had small lesions, which were challenging to accurately segment without proper adjustment of patch size and use of a patch restriction strategy. With our final strategy,

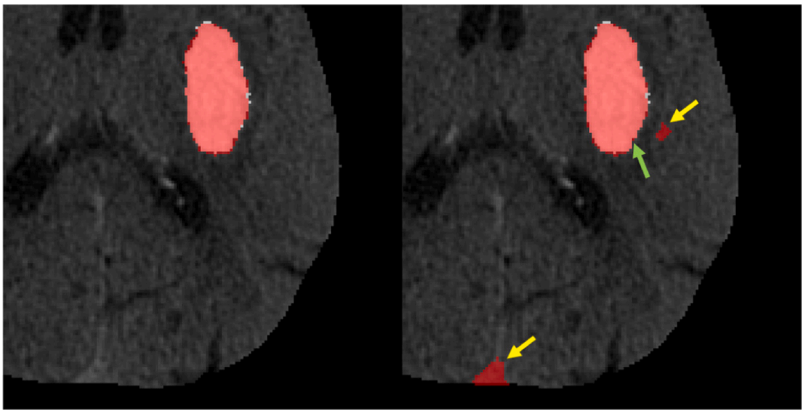

Fig. 9. Example of miss-segmentation. Even though ROI helped to refine contours (green arrow), voxels not related to the lesion were also segmented (yellow arrows).

almost all of them achieved the DSC > 0.8, except for one case (pt030) with DSC $= 0.742$ due to the beam hardening artifact presented on this CT image. However, we consider it as a big improvement compared to the initial experiments, where the DSC for this case could be as low as 0.08. In total, 85% of all cases reached the threshold of DSC > 0.8 and 33% of all cases have a DSC > 0.9.

5. Conclusions

In this work we proposed a deep learning method for hemorrhagic stroke lesions segmentation in CT images. The proposed approach, based on a patch based 3D U-Net architecture with integration of squeeze-and-excitation blocks, was tested on a clinical dataset of 76 cases. The obtained results were qualitatively and quantitatively evaluated on the whole dataset using a 5-fold cross-validation strategy. Our experiments demonstrated that such an architecture significantly improved segmentation results compared to the baseline 3D U-Net approach. Moreover, we showed also that data preparation is a key step to obtain good segmentation results. By using a restrictive balanced sampling technique, we successfully tackled the class imbalance problem as well as the problem of intra-ventricular hemorrhage present in some of the cases in the dataset. In addition, the inclusion of the symmetric modality as additional input channel helped to improve the results, obtaining a mean DSC segmentation result of 0.862 ± 0.074 with an approximated segmentation time of 17 s per patient.

Conflict of interest statement

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

Declaration of Competing Interest

The authors report no declarations of interest.

Acknowledgments

This work has been supported by Retos de Investigación DPI2017-86696-R from the Ministerio de Ciencia, Innovación y Universidades and by Universitat de Girona PONTUdG2020/09. Valeriia Abramova holds an EACEA Erasmus+ grant for the master in Medical Imaging and Applications (MAIA), and Albert Clèrigues a FPI grant from the Ministerio de Ciencia, Innovación y Universidades (PRE2018-083507). The authors gratefully acknowledge the support of the NVIDIA Corporation with their donation of the GPU used in this research.

References

Arab, A., Chinda, B., Medvedev, G., Siu, W., Guo, H., Gu, T., Moreno, S., Hamarneh, G., Ester, M., Song, X., 2020. A fast and fully-automated deep-learning approach for accurate hemorrhage segmentation and volume quantification in non-contrast whole-head CT. Sci. Rep. 10, 19389. https://doi.org/10.1038/s41598-020-76459-7.

Bardera, A., Boada, I., Feixas, M., et al., 2009. Semi-automated method for brain hematoma and edema quantification using computed tomography. Comput. Med. Imaging Graphics 33, 304–311.

Bernal, J., Kushibar, K., Cabezas, M., Valverde, S., Oliver, A., Lladć, X., 2019. Quantitative analysis of patch-based fully convolutional neural networks for tissue segmentation on brain magnetic resonance imaging. IEEE Access 7, 89986–90002.

Chang, P., Kuoy, E., Grinband, J., et al., 2018. Hybrid 3d/2d convolutional neural network for hemorrhage evaluation on head ct. Am. J. Neuroradiol.

Chinda, B., Medvedev, G., Siu, W., Ester, M., Arab, A., Gu, T., Moreno, S., D'Arcy, R.C.N., Song, X., 2018. Automation of ct-based haemorrhagic stroke assessment for improved clinical outcomes: study protocol and design. BMJ Open 8. https://doi.org/10.1136/bmjopen-2017-020260.

Cho, J., Park, K.S., Karki, M., et al., 2019. Improving sensitivity on identification and delineation of intracranial hemorrhage lesion using cascaded deep learning models. J. Digital Imaging 32.

Clèrigues, A., Valverde, S., Bernal, J., et al., 2019. Acute ischemic stroke lesion core segmentation in ct perfusion images using fully convolutional neural networks. Comput. Biol. Med. 115, 103487.

Cosić, D., Lončarić, S., 1997. Computer system for quantitative: analysis of ich from ct head images. Annual International Conference of the IEEE Engineering in Medicine and Biology Society vol. 2 553–556.

De Oliveira Manoel, A., Goffi, A., Zampieri, F., Turkel-Parrella, D., Duggal, A., Marotta, T., Abrahamson, S., 2016. The critical care management of spontaneous intracranial hemorrhage: a contemporary review. Critical Care 20. https://doi.org/10.1186/s13054-016-1432-0.

Hakimi, R., Garg, A., 2016. Imaging of hemorrhagic stroke. CONTINUUM: Lifelong Learning in Neurology 22, 1424–1450. https://doi.org/10.1212/CON.0000000000000377.

Heit, J.J., Iv, M., Wintermarkl, M., 2017. Imaging of intracranial hemorrhage. J. Stroke 19, 11–27.

Hemphill, J., Bonovich, D., Besmertis, L., Manley, G., Johnston, s.C., 2001. The ich score a simple, reliable grading scale for intracerebral hemorrhage. Stroke 32, 891–897. https://doi.org/10.1161/01.STR.32.4.891.

Hinson, H., Hanley, D., Ziai, W., 2010. Management of intraventricular hemorrhage. Curr. Neurol. Neurosci. Rep. 10, 73–82.

Hssayeni, M., Al-Janabi, M., Salman, A., et al., 2020. Intracranial hemorrhage segmentation using a deep convolutional model. Data 5, 14.

Hu, K., Chen, K., He, X., et al., 2020. Automatic segmentation of intracerebral hemorrhage in ct images using encoder-decoder convolutional neural network. Inform. Process. Manage. 57, 102352.

Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. NeuroImage 17, 825–841.

Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. Med. Image Anal. 5, 143–156.

Kidwell, C., Wintermark, M., 2008. Imaging of intracranial haemorrhage. Lancet Neurol. 7, 256–267.

Kuang, Z., Deng, X., Yu, L., et al., 2020. ψ. Comput. Methods Progr. Biomed. 194, 105546.

Kuo, W., Häne, C., Mukherjee, P., et al., 2019. Expert-level detection of acute intracranial hemorrhage on head computed tomography using deep learning. Proc. Natl. Acad. Sci. 116, 22737–22745.

Kushibar, K., Valverde, S., González-Villà, S., et al., 2018. Automated sub-cortical brain structure segmentation combining spatial and deep convolutional features. Med. Image Anal. 48, 177–186.

Kwon, D., Ahn, J., Kim, J., et al., 2019. Siamese u-net with healthy template for accurate segmentation of intracranial hemorrhage. Medical Image Computing and Computer Assisted Intervention - MICCAI 2019 848–855.

Lin, T., Goyal, P., Girshick, R., He, K., Dollér, P., 2017. Focal loss for dense object detection. 2017 IEEE International Conference on Computer Vision (ICCV) 2999–3007.

Lončarić, S., Dhawan, A., Broderick, J., Brott, T., 1995. 3-d image analysis of intracerebral brain hemorrhage from digitized ct films. Comput. Methods Programs Biomed. 46, 207–216.

Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation. 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 3431–3440.

Matsuo, R., Yamaguchi, Y., Matsushita, T., et al., 2017. Association between onset-to-door time and clinical outcomes after ischemic stroke. Stroke 48.

Mozaffarian, D., Benjamin, E., Go, A., et al., 2016. Heart disease and stroke statistics-2016 update a report from the American heart association. Circulation 133, e38–e48.

Paszke, A., Gross, S., Chintala, S., et al., 2017. Automatic differentiation in pytorch. NIPS-W.

Patel, A., Schreuder, F., Klijn, C., et al., 2019. Intracerebral haemorrhage segmentation in non-contrast ct. Sci. Rep. 9. https://doi.org/10.1038/s41598-019-54491-6.

Perez, N., Valdes, J., Guevara Lopez, M.A., Rodríguez, L., Molina, J., 2007. Set of methods for spontaneous ICH segmentation and tracking from ct head images 212–220.

Puig, J., Blasco, G., Daunis-i Estadella, P., et al., 2017. High-permeability region size on perfusion ct predicts hemorrhagic transformation after intravenous thrombolysis in stroke. PLOS ONE 12, e0188238.

Roth, G.A., Abate, D., Abate, K.H., et al., 2018. Global, regional, and national age-sex-specific mortality for 282 causes of death in 195 countries and territories, 1980-2017: a systematic analysis for the global burden of disease study 2017. The Lancet 392, 1736–1788.

Shahangian, B., Pourghassem, H., 2015. Automatic brain hemorrhage segmentation and classification algorithm based on weighted grayscale histogram feature in a hierarchical classification structure. Biocybernet. Biomed. Eng. 36.

Sharrock, M., Mould, W., Ali, H., Hildreth, M., Awad, I., Hanley, D., Muschelli, J., 2020. 3d deep neural network segmentation of intracerebral hemorrhage: Development and validation for clinical trials. Neuroinformatics. https://doi.org/10.1007/s12021-020-09493-5.

Singh, S., Ker, J., Bai, Y., et al., 2019. Image thresholding improves 3-dimensional convolutional neural network diagnosis of different acute brain hemorrhages on computed tomography scans. Sensors 19, 2167.

Soun, J., Chow, D., Nagamine, M., Takhtawala, R., Filippi, C., Yu, W., Chang, P., 2020. Artificial intelligence and acute stroke imaging. Am. J. Neuroradiol. https://doi.org/10.3174/ajnr.A6883.

Wang, Y., Liu, H., Liu, Y., Liu, W., 2018. Deep learning framework for hemorrhagic stroke segmentation and detection. BIBE 2018; International Conference on Biological Information and Biomedical Engineering 1–6.

Wismüller, A., Stockmaster, L., 2020. A prospective randomized clinical trial for measuring radiology study reporting time on Artificial Intelligence-based detection

of intracranial hemorrhage in emergent care head CT. In: Krol, A., Gimi, B.S. (Eds.), Medical Imaging 2020: Biomedical Applications in Molecular, Structural, and Functional Imaging, International Society for Optics and Photonics. SPIE, pp. 144–150. https://doi.org/10.1117/12.2552400.

Woo, I., Lee, A., Jung, S., et al., 2019. Fully automatic segmentation of acute ischemic lesions on diffusion-weighted imaging using convolutional neural networks: Comparison with conventional algorithms. Korean J. Radiol. 20, 1275.

Yao, H., Williamson, C., Gryak, J., Najarian, K., 2020. Automated hematoma segmentation and outcome prediction for patients with traumatic brain injury. Artif. Intel. Med. 107, 101910.

Zeiler, M.D., 2012. Adadelta: An adaptive learning rate method. arXiv:1212.5701.