



Combining unsupervised and supervised learning for predicting the final stroke lesion



Adriano Pinto^{a,b,*}, Sérgio Pereira^{a,b}, Raphael Meier^d, Roland Wiest^d, Victor Alves^b,
Mauricio Reyes^c, Carlos A. Silva^a

^a Center MEMS of University of Minho, Campus of Azurém, Guimarães 4800-058 Portugal

^b Center Algoritmi, University of Minho, Braga, Portugal

^c Healthcare Imaging A.L., Insel Data Science Center, Bern University Hospital, Switzerland

^d Support Center for Advanced Neuroimaging, University Institute for Diagnostic and Interventional Neuroradiology, Bern University Hospital, Switzerland

ARTICLE INFO

Article history:

Received 22 January 2020

Revised 9 October 2020

Accepted 22 October 2020

Available online 24 December 2020

Keywords:

Deep learning

Image prediction

Magnetic resonance imaging

Stroke

ABSTRACT

Predicting the final ischaemic stroke lesion provides crucial information regarding the volume of salvageable hypoperfused tissue, which helps physicians in the difficult decision-making process of treatment planning and intervention. Treatment selection is influenced by clinical diagnosis, which requires delineating the stroke lesion, as well as characterising cerebral blood flow dynamics using neuroimaging acquisitions. Nonetheless, predicting the final stroke lesion is an intricate task, due to the variability in lesion size, shape, location and the underlying cerebral haemodynamic processes that occur after the ischaemic stroke takes place. Moreover, since elapsed time between stroke and treatment is related to the loss of brain tissue, assessing and predicting the final stroke lesion needs to be performed in a short period of time, which makes the task even more complex. Therefore, there is a need for automatic methods that predict the final stroke lesion and support physicians in the treatment decision process. We propose a fully automatic deep learning method based on unsupervised and supervised learning to predict the final stroke lesion after 90 days. Our aim is to predict the final stroke lesion location and extent, taking into account the underlying cerebral blood flow dynamics that can influence the prediction. To achieve this, we propose a two-branch Restricted Boltzmann Machine, which provides specialized data-driven features from different sets of standard parametric Magnetic Resonance Imaging maps. These data-driven feature maps are then combined with the parametric Magnetic Resonance Imaging maps, and fed to a Convolutional and Recurrent Neural Network architecture. We evaluated our proposal on the publicly available ISLES 2017 testing dataset, reaching a Dice score of 0.38, Hausdorff Distance of 29.21 mm, and Average Symmetric Surface Distance of 5.52 mm.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Stroke is the second leading cause of death worldwide (World Health Organization et al., 2014), being classified in two types: ischaemic and haemorrhagic (Grysiewicz et al., 2008). Ischaemic stroke is the most common type, resulting from an occlusion of a vessel, which can be caused by thrombolysis, haemodynamic factors, or embolic causes (Grysiewicz et al., 2008). Due to vessel occlusion, the insufficient supply of oxygenated blood to brain cells leads to hypoperfused brain tissue, triggering cellular mechanisms to preserve the integrity of the cell. The hypoperfused

area consists of tissue at risk that can be salvaged, being designated penumbra. As time passes, in the absence of flow restoration or sufficient collateral blood flow supply, the hypoperfused tissue eventually reaches a non-salvageable state designated core or infarct tissue (Memezawa et al., 1992).

Diagnosis and treatment of ischaemic stroke relies on neuroimaging acquisitions, where Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) are the preferred imaging modalities (Gonzalez et al., 2007). CT imaging remains the most used acquisition due to its rapidity and availability (Gonzalez et al., 2007). However, multi-parametric MRI provides a higher sensitivity in detecting early ischaemic stroke and assessing the penumbra region (Gonzalez et al., 2007). Treatment consists in restoring tissue perfusion levels, also known as reperfusion, by performing mechanical thrombectomy or thrombolysis. Since ischaemic stroke is a dynamic process that evolves over time, the treatment

* Corresponding author at: Department of Industrial Electronics, Campus Azurém, Guimarães, Portugal

E-mail addresses: id6376@alunos.uminho.pt (A. Pinto), csilva@dei.uminho.pt (C.A. Silva).

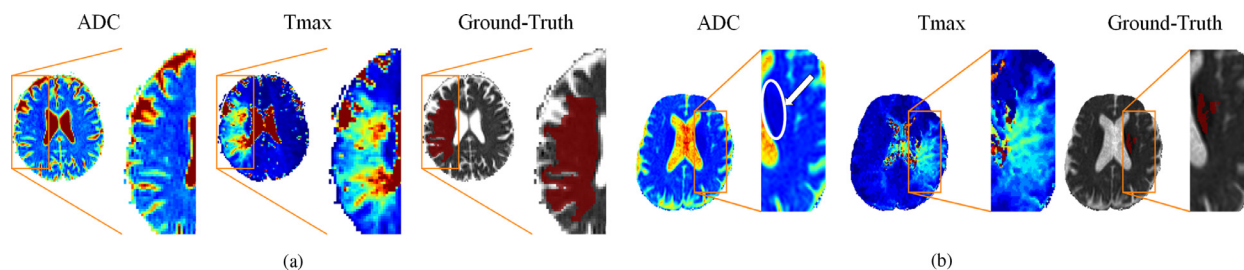


Fig. 1. ADC and Tmax parametric maps of two patient cases from ISLES 2017 training set, and the final lesion delineated at a 90-day follow-up, overlapped with the onset ADC: patient 0036 (Fig. 1a) with an unsuccessful reperfusion, and patient 0006 (Fig. 1b), where the clinical intervention was successful.

is only possible up to 24 h, where viable neurones still persist (El Tawil and Muir, 2017; Zivelonghi and Tamburini, 2018). So, expert physicians must evaluate the benefits and risks of mechanical thrombectomy before an intervention, since it may cause haemorrhage, vascular injury, and other complications (Powers et al., 2018). If performed, the success of the intervention is assessed radiologically via angiography imaging and scored by a qualitative expert-generated scale designated the standardized Thrombolysis in Cerebral Infarction (TICI) scale (Higashida et al., 2003). During the decision-making process, the physician needs to assess the nature and location of the lesion alongside pathophysiological factors such as age, presence of comorbidities, and collateral circulation (Liebeskind, 2003). The latter is of utmost importance in ischaemic stroke. The presence of collateral circulation, where a secondary network of vessels is responsible for granting cerebral blood flow to the lesioned tissue, increases the chances of a successful reperfusion (Liebeskind, 2003). Asserting the potential efficacy of treatment can be time-consuming and prone to inter- and intra-variability among physicians, which is further potentiated when performed in a clinical emergency environment (Coutts et al., 2003). Moreover, since time is critical, MRI acquisitions are optimized for speed, which is accomplished by reducing the resolution (González et al., 2011), making the prediction of the final stroke lesion an intricate task. Thus, automatic prediction of a stroke lesion at a given time since stroke has a great potential to guide physicians in this time-critical decision-making process.

We propose a novel automatic method based on unsupervised and supervised deep learning. We utilize Restricted Boltzmann Machines (RBMs) to jointly characterise the lesion and blood flow information through a two-pathway architecture, trained with two subsets of standard parametric MRI maps. One subset encompasses the Time-To-Peak (TTP), Mean Transit Time (MTT), Time-to-Maximum (Tmax), and Apparent Diffusion Coefficient (ADC). The second set contains the ADC, the relative Cerebral Blood Volume (rCBV), and the relative Cerebral Blood Flow (rCBF). In a second stage, the feature maps computed by the RBMs are combined with the standard parametric MRI maps to form the input of a supervised deep learning architecture composed by Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). The proposed architecture was evaluated using the publicly-available ISLES 2017 dataset.

1.1. On the complexity of predicting the final infarct stroke lesion

In acute ischaemic stroke, the clinical evaluation of the standard parametric maps (e.g. ADC and Tmax) can identify infarct tissue and tissue that will infarct in the absence of therapeutic intervention. In this analysis, the infarct tissue, is identified by the hypointense regions of the ADC map, which characterize tissue with limited diffusion (Butcher and Emery, 2010a). Hypoperfused tissue, i.e. tissue that will infarct, is identified by hyperintense regions of the Tmax map, indicating an increased arrival time of contrast agent (Butcher and Emery, 2010b). However, to correctly predict

the final ischaemic stroke lesion, besides considering the complex time-evolving transformation of hypoperfused tissue to infarcted tissue, it is also necessary to appraise the impact of the clinical intervention, thrombectomy, on the underlying brain perfusion and diffusion.

A successful thrombectomy should restore the perfusion levels, recovering the hypoperfused tissue. However, several factors may affect the reperfusion, limiting the degree of success of the intervention. To better understand the nuances of the clinical intervention, consider the cases presented in Fig. 1. In the first case, Fig. 1a, the ADC does not present any hypointense region, so no infarct tissue may be identified, and we should expect a complete recovery of the hypoperfused tissue indicated by Tmax; however, the follow-up delineation obtained after thrombectomy presents a large final lesion, which is explained by an unsuccessful intervention. In the second case, Fig. 1b, we observe a final infarct lesion that is smaller than the hypointense region present in the ADC (Fig. 1b arrow). This indicates reversible diffusion restriction, which is a rare case (Labeyrie et al., 2012) and was only possible to identify by a follow-up T2-weighted acquisition. So, an automatic method for predicting the final stroke lesion has not only to capture the time-evolving process of diffusion and perfusion, but also to consider directly or indirectly the degree of success of the thrombectomy, which may condition the final lesion either to be confined to the hypointense region of the ADC map, or to grow to brain tissue areas that are hyperintense in the Tmax. Due to the time-evolving process of diffusion and perfusion, the complexity of predicting the lesion will aggravate as we move from a target window of some days to several months.

The complexity of the evaluation process may be also observed in the inter-rater agreement of expert radiologists in ISLES 2017 dataset, which obtained a Dice score of 0.58 ± 0.20 on delineating the lesion using a 90-day follow-up T2-weighted acquisition (Winzeck et al., 2018).

1.2. Previous work

Contrary to stroke lesion segmentation, where several methods have already been proposed (Rekik et al., 2012; Maier et al., 2017), the complexity of predicting the final stroke lesion has only recently attracted attention in the medical imaging community. For predicting the final stroke lesion several methods have been already proposed based on multivariate linear regression models (Scalzo et al., 2012; Rose et al., 2001; Kemmling et al., 2015), decision trees (McKinley et al., 2017; Bauer et al., 2014), and CNNs (Choi et al., 2016). Furthermore, with the release of Ischaemic Stroke LESion Segmentation (ISLES) Challenge in 2016 and 2017, new methods have been proposed. These aim to predict at a 90-day time-window.

Rose et al. (2001) proposed a two-stage method based on parametric perfusion and diffusion MRI maps. On the first stage, the method defines a region of interest (ROI) based on the intensity signal of the standard parametric maps, the MTT, Cerebral Blood

Flow (CBF), Cerebral Blood Volume (CBV), and Diffusion-Weighted Imaging (DWI). The second stage performs stroke tissue prediction using Gaussian mixture models trained in different sets of parametric maps. Bauer et al. (2014) used Random Forests to segment or predict the final stroke lesion depending on whether acute stroke imaging or three-month follow-up imaging was available, respectively. McKinley et al. (2017) also used a two-stage classification approach as in Rose et al. (2001) for lesion characterisation and lesion prediction, where each stage consists of two sets of Random Forests (RFs) classifiers. The first stage aims to define a ROI that encompasses the hypoperfused region. In the first set, each classifier is trained with features extracted from different sets of MRI parametric maps. Having defined the location and extension of the lesion, a second set of two RFs performs stroke tissue prediction. Such classifiers were trained on different sets of patients, stratified by the TICI score. One classifier is trained in patients with unsuccessful reperfusion interventions, whereas a second classifier is trained in patients with successful reperfusion. The final prediction is obtained by combining the results of both classifiers, using a logistic regression model. Scalzo et al. (2012) proposed a framework for stroke tissue prediction, which characterises the state of the lesion four days after clinical intervention (thrombectomy). From the Fluid Attenuation Inversion Recovery (FLAIR) MRI sequence, ADC and Tmax MRI maps, the method applies a regression model that learns the behaviour of neighbouring voxels within a cuboid. Kemmling et al. (2015) proposed a multi-modality approach based on CT and MRI maps with non-imaging clinical meta-data, namely the TICI score and the time to treatment of each patient, to perform stroke tissue prediction.

In another line of research, authors have investigated the use of deep learning (Choi et al., 2016; Nielsen et al., 2018; Robben et al., 2020) for stroke tissue prediction. Choi et al. (2016), the winner approach at ISLES 2016 Challenge, proposed an ensemble of twelve CNN architectures, grouped into two sets of networks. The first group comprehends four 3D U-Nets (Ronneberger et al., 2015) performing voxel-wise tissue prediction. The second group of networks uses two-pathway Fully Connected Networks (FCNs) performing two types of patch-wise classification. One set of FCNs classifies a patch as lesion if it includes any lesion voxel. The other set of FCNs classifies a patch as lesion if the central voxel is a lesion. After merging the two pathway FCN, the method incorporates meta-data by adding a dense layer of clinical predictors merged with the imaging output of each network. The final stroke lesion prediction results from a weighted merging of all models. Mok and Chung (2017) applied deep adversarial training for stroke tissue prediction in an ensemble of U-Nets. Monteiro and Oliveira (2017) proposed a method based on the V-Net architecture (Milletari et al., 2016). The training was conducted with a custom loss function that applies a weighted sum between Dice score and cross entropy. Lucas and Heinrich (2017) proposed the use of a U-Net architecture, which combines patches from the MRI maps in the same slice, with patches from 3 neighbouring slices and 2 hemispheric flips. In the expanding path of the U-Net, each level computes a Dice loss for the healthy tissue and for the final lesion, after the softmax activation. Afterwards, all losses are summed up, having the loss of the lesion and healthy tissue weighted according to a prior probability (Winzeck et al., 2018). Robben and Suetens (2017) employed a CNN-based architecture inspired by Kamnitsas et al. (2017). The authors proposed to combine the MRI inputs with clinical meta-data, before feeding them to each branch of a two-pathway 3D network. In the first branch the input is kept with the original resolution, while in the second branch the input resolution was lowered by a factor of 3. The output of each branch is transformed to the same scale and merged by two fully connected layers. The network is trained with four different sets of hyper-parameters.

These four networks are used as an ensemble, whose prediction is obtained by averaging the output of each one. Similarly, Niu et al. (2018) used multiple scales of overlapping 3D patches to capture local and global spatial information. In the review paper of Winzeck et al. (2018), Rivera et al. also built on the work of Kamnitsas et al. (2017) and Milletari et al. (2016), by proposing a scheme to extract different patch resolutions, independent of each other, that are fed into four different paths. Afterwards, a fully connected layer combines all the outputs to perform stroke tissue prediction. Pisov et al. (2017) employed an ensemble strategy by combining different CNN-based architectures to overcome the strong anisotropy of the data. As summarized by Winzeck et al. (2018), the work of Yoon et al. proposed a two-stage gated CNN. In a first stage, the authors perform lesion detection and delineation. Afterwards, based on the probability maps of the first stage, a second CNN architecture processes the regions where the probability maps of healthy tissue and lesion are close to each other. Pinto et al. (2018b) made use of temporal perfusion imaging, the Dynamic Susceptibility Contrast-MRI, in a U-Net architecture. This architecture aims to temporally process and extract deep features, which are then combined with a second feature step of another U-Net network, which was trained on the standard parametric maps. Using a large CT dataset, Robben et al. (2020) predicted the final infarct stroke lesion with a temporal window ranging from 24 h to 5 days. The authors considered spatio-temporal CT perfusion as input to a deep neural network inspired in the architecture proposed by Kamnitsas et al. (2017). Additionally, the model combines CT neuroimaging with clinical meta-data. Nielsen et al. (2018) proposed a method based on the SegNet architecture (Badrinarayanan et al., 2015), predicting on a 30-day follow-up acquisition based on a private dataset.

Principal and collateral blood flow has been considered either directly by modelling the temporal perfusion imaging (Pinto et al., 2018b), or indirectly by perfusion and diffusion parametric maps (Choi et al., 2016; Maier et al., 2017; Scalzo et al., 2012), or through clinical information that characterises the success of the revascularization (McKinley et al., 2017). We hypothesize that modelling the haemodynamics of the brain when artery occlusion occurs can be beneficial for predicting the final stroke lesion. So, in this work, we investigate the representation of the haemodynamics through an unsupervised learning model. Contrary to previous approaches, we propose grouping the input maps according to their subjacent physical meaning and encoding each group separately with an RBM. As groups, we investigated the time-resolved perfusion maps (Tmax, TTP, MTT), and the blood-flow-dynamic related maps (rCBF, rCBV) (Butcher and Emery, 2010a; 2010b). Our proposal of combining features obtained unsupervisedly and supervisedly was motivated by the knowledge that unsupervised models learn structural features of the original image, while the supervised models learn features conditioned on the label, so there is potential for obtaining richer and more discriminative features by joining both types of models.

1.3. Contributions

This work presents an automatic approach for predicting the final stroke lesion, using onset neuroimaging data. The main contributions are:

- The use of unsupervised models for extracting structural features of time-resolved perfusion and blood-flow-dynamic related MRI maps for predicting stroke lesion.
- The use of local and long spatial context provided by gated recurrent neural networks for relating structural features and image information when learning features conditioned on the label in a supervised model.

- The proposal of a competitive system which outperforms state-of-the-art methods to predict the final infarct stroke lesion, in ISLES 2017 Challenge dataset.

The remainder of the paper is organized as follows. Section 2 describes the fundamental components of the proposed method. Section 3 describes the dataset, the evaluation procedure and the setup. The results and the discussion are addressed in Section 4. Finally, in Section 5 we present the main conclusions.

2. Methods

In this work, predicting the final infarct stroke lesion consists of delineating the lesion’s spatial extent at a 90-day follow-up time-point, using multi-parametric MRI imaging, namely the ADC, MTT, TTP, Tmax, rCBF, and rCBV, which are acquired at the onset time-point. The architecture of the proposed system and its main components are described in the following subsections.

2.1. Architecture

The overall architecture of the proposed method can be divided into two functional blocks, as shown in Fig. 2.

The first functional block performs unsupervised representation learning using two unsupervised models, namely RBMs. This unsupervised block provides new features that represent structural information that complements the standard parametric MRI maps, enhancing the capacity of our model to predict the final infarct lesion. In our approach, we aim to model the clinical procedure, which first locates and delineates the lesion at current time, and then considers the blood flow haemodynamic that might influence the final stroke lesion prediction. This procedure is encoded in our two-path RBM. The first RBM is responsible for capturing information on lesion location and extension, referred to as the RBM_{Lesion} . The second RBM, RBM_{Haemo} , aims to capture blood flow haemodynamics information (e.g. collateral circulation), which has been identified as a key factor by physicians when assessing stroke final infarct lesion in clinical reports (Berkhemer et al., 2016; Menon et al., 2015). On one hand, to locate the onset ischaemic stroke lesion, the RBM_{Lesion} considers standard parametric maps that characterise the arrival times and mean transit times of the contrast agent. In the presence of an ischaemic lesion, the occluded vessel can decrease or interrupt the normal brain perfusion, translating into hyperintense regions on time-related parametric maps (Butcher and Emery, 2010b). On the other hand, the RBM_{Haemo} considers standard parametric maps that characterise the amount of blood being delivered in unit of time, which correlates to the cerebral blood flow haemodynamics (Butcher and

Emery, 2010b). Thus, the RBM_{Lesion} considers the MTT, TTP and Tmax perfusion maps, while the RBM_{Haemo} the rCBV and rCBF perfusion maps. Regarding the ADC standard diffusion map, it is present in both RBM_{Lesion} and RBM_{Haemo} , since it provides higher brain structural information and allows the identification of tissue that is already infarcted. This separation of the input imaging allows the RBM to learn specific feature sets, which may enable the method to analyse difficult cases where information concerning the blood flow can have a favourable impact on the lesion prediction.

The second functional block consists of a deep learning architecture that comprehends 2D convolutional blocks in a U-Net-based structure, alongside recurrent blocks. As imaging input data, we combine the standard parametric maps with feature maps from each RBM, totalling 18 input feature maps.

2.2. Restricted Boltzmann machines

The RBM is an undirected graphical model constituted by two layers of nodes: a visible layer and a hidden layer (Rumelhart and McClelland, 1986). Each node has a weighted connection to all nodes in the other layer (Rumelhart and McClelland, 1986). However, there are no connections among nodes of the same layer. Originally, Rumelhart and McClelland (1986) proposed the RBM to learn from binary data on both layers. However, this does not represent well continuous real-valued input data, which is the case of MRI data. Therefore, we model the visible nodes as linear units with independent Gaussian noise. The hidden nodes are modelled as Noisy Rectifier Linear Units (NReLU), since they have been reported to be suitable for feature extraction (Hinton, 2012). This kind of RBM was previously used in segmentation tasks, such as in Pereira et al. (2019). Mapping the input data into a feature vector is performed through the interaction of states between the visible and hidden units, which is learned by minimizing an energy function.

The complete pipeline of the unsupervised block is shown in Fig. 3 and detailed in Section 3.4. The RBM_{Lesion} and RBM_{Haemo} function as feature generators that output two complementary sets of feature maps \mathcal{N}_1 and \mathcal{N}_2 . These features characterise the structure of the images; however, we are interested only on the most distinctive details. So, after training the RBMs, we perform feature selection to reduce the generated feature space, obtaining smaller but representative feature sets \mathcal{M}_1 and \mathcal{M}_2 , such that $|\mathcal{M}_i| \ll |\mathcal{N}_i|$, for $i \in [1, 2]$, where the operator $|\cdot|$ denotes the cardinality of a set. In the feature selection, we would like to select the features from the RBM that encodes the MRI maps, but also that correlates with the stroke prediction. Since the RBM is an unsupervised method, we compute the Normalized Mutual Information to quantify the statistical dependence between each generated feature and the

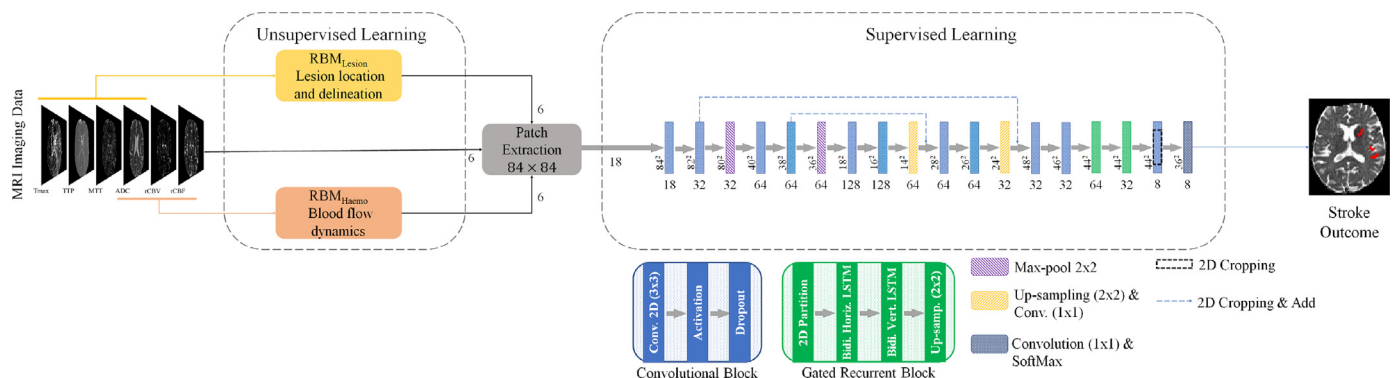


Fig. 2. Overview of the proposed method for predicting the final stroke lesion. In the supervised learning block, the input data dimensions are defined for each operation.

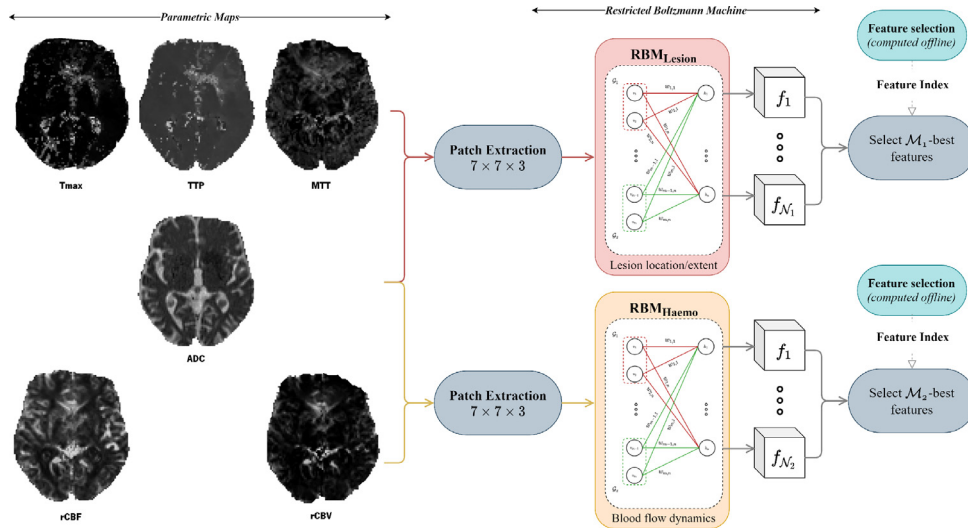


Fig. 3. Overview of the proposed unsupervised learning block. For each RBM of the unsupervised learning block, the selected features were $\mathcal{M}_1 = \mathcal{M}_2 = 6$.

respective input MRI map, as defined by Eq. (1) (Vinh et al., 2010):

$$NMI_{sum}(MRI_x, Feat_y) = 2 \frac{MI(MRI_x, Feat_y)}{H(MRI_x) + H(Feat_y)}, \quad (1)$$

where $MI(\cdot)$ is the mutual information between an MRI parametric map, MRI_x , and an output feature, $Feat_y$; $H(\cdot)$ defines the entropy of a map, namely, MRI_x and $Feat_y$. To relate the features of the RBM with the class label, we could use a classifier supervisedly trained. Since the neural network is trained iteratively, we use a RF classifier trained with the Mean Decrease Impurity (MDI) as a surrogate to make the feature selection tractable. After, we compute the MI_{RBM} and MDI_{RF} , normalize the MI_{RBM} by the maximum value, add both ranks and sort decreasingly. The best set will be the first \mathcal{M}_i features. Our selection method was inspired on the work of Pereira et al. (2018); however, their method cannot be directly applied, since it would generate too many features for our problem.

2.3. Convolutional and recurrent neural networks

Our supervised functional block is based on the U-Net architecture as proposed by Ronneberger et al. (2015). The input of the U-Net considers the concatenation of standard parametric maps with the sets of feature maps extracted from the unsupervised block. In the first level of our encoder architecture we use four 2D convolutional blocks with kernel size of 3×3 and 32 channels. Afterwards, the output of the final convolutional block is down-sampled by a factor of 2, starting the second encoding level formed by two convolutional blocks with equal kernel size but doubling the number of feature maps. The third level of encoding follows the same pattern. The decoder level mimics the encoder counterpart. As in Ronneberger et al. (2015) we only used long skip connections among encoder and decoder levels. These encoder-decoder deep CNNs provide high levels of abstraction from the input data, increasing the global notion of context as the network grows deeper. However, it comes at a cost of a high receptive field (Zeiler and Fergus, 2014). Thus, we used a 2D architecture in the plane with the highest resolution, since the acquisition resolution is anisotropic in the dataset. Also, in the end of the decoding path we expanded our learning block with Gated RNNs. Due to their nature, Gated RNNs can capture short- and long-term spatial relations, by retaining information from previous nodes encoded in the time-steps. Hence, Gated-RNNs consider information from all previous nodes when analysing the current one.

This property, when applied to imaging data, allows considering intra-slice contextual dependencies. In our work, we used a particular Gated-RNN, namely the Long-Short Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997). However, the LSTM was intrinsically developed to process 1D data (Hochreiter and Schmidhuber, 1997) (e.g. time-series). To be applicable to 2D data, we developed an online 2D Partition layer that transforms a grid-structure input (e.g. an image) into a one-dimensional sequence. Inspired by Visin et al. (2016), the 2D Partition layer was predefined with a neighbourhood of 2×2 , where each time-step is characterised by a feature space of four voxels. After, two Bidirectional LSTM layers are employed along the left-right and frontal-dorsal directions followed by an up-sampling layer. These four layers, referred as the Gated Recurrent block, are shown in Fig. 2. In our supervised functional block, two Gated Recurrent blocks were used, where the Bidirectional LSTMs have 64 and 32 hidden layers, respectively. The impact of the main components is evaluated in an ablation study in the experiments.

3. Experimental setup

We evaluated the proposed approach on the publicly available ISLES 2017 dataset and on a private dataset. ISLES has an online benchmark platform (Kistler et al., 2013) that performs automatic evaluation (SMIR Online Platform, 2017). In this section we describe the dataset, the training and evaluation, and the main hyper-parameters of our method.

3.1. Data

ISLES 2017 dataset encompasses 75 ischaemic stroke patients, which are separated into two sets: training ($n = 43$) and testing ($n = 32$). Both sets have patients that underwent mechanical thrombectomy. Each patient is characterised by six 3D parametric MRI maps: diffusion ADC map, perfusion rCBF, rCBV, TTP, MTT and Tmax maps. In addition to the standard parametric maps, each case is also characterised by a manual delineation of the lesion. This refers to the 90-day stroke lesion delineated with access to the follow-up T2-weighted acquisition. However, the manual delineation is only available for the training set, while the follow-up T2-weighted imaging is not disclosed for any set. All parametric MRI maps are already co-registered and skull-stripped (Winzeck et al., 2018). Fig. 4 (top row) shows an example of MRI maps, alongside the manual lesion delineation, the Ground Truth (GT), of a patient.

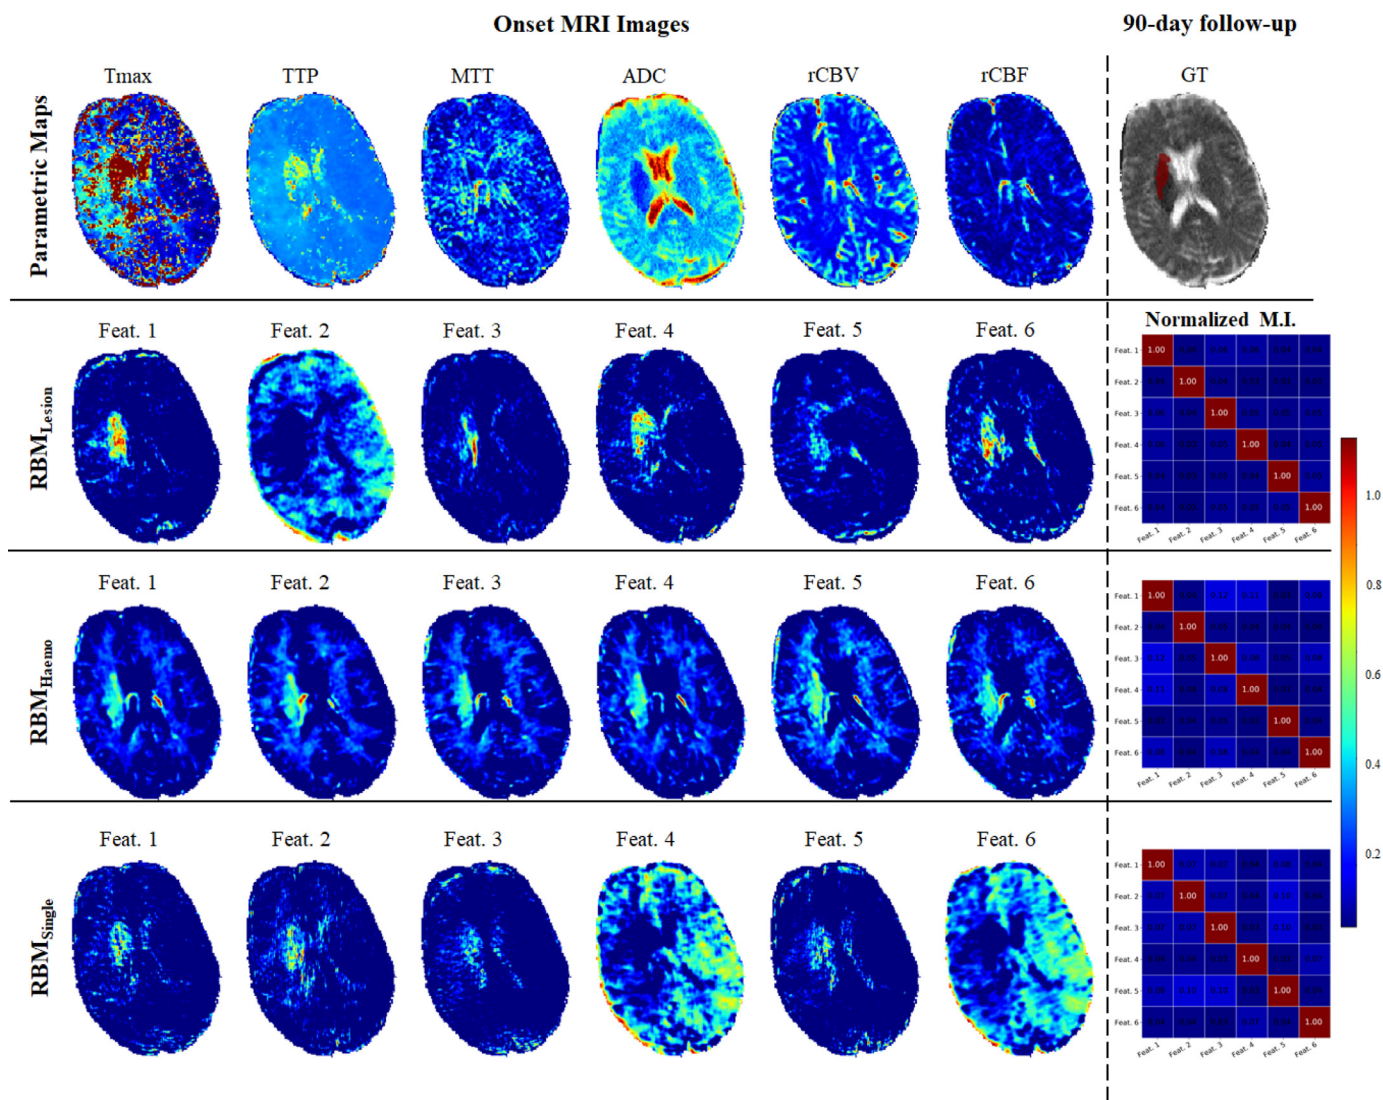


Fig. 4. Onset parametric maps of patient case 0011 in ISLES 2017 training set, alongside the final stroke lesion, at a 90-day follow-up, over the onset ADC map. The subsequent rows show the RBM features selected from the RBM_{Lesion} , RBM_{Haemo} and RBM_{Single} , respectively. The last column shows the normalized mutual information, across whole dataset, among features of the same RBM.

The private dataset considers 23 acute ischaemic stroke patients that underwent clinical therapy, acquired at Bern University Hospital in Switzerland. As in ISLES 2017, each patient is characterized by the same six parametric maps, being the final lesion manually delineated at 90-day follow-up T2. The parametric maps were co-registered followed by skull-stripping with FSL BET2 on the co-registered follow-up T2 image (Jenkinson et al., 2012).

3.2. Evaluation metrics

We evaluated our proposal with five metrics, which are the same ones computed by the online ISLES 2017 benchmark platform: Dice Similarity Score, Hausdorff Distance (HD), Average Symmetric Surface Distance (ASSD), Precision, and Recall (Kistler et al., 2013).

Dice score measures the spatial overlap between two volumes. HD corresponds to the highest distance between surface points of different volumes, which characterise spatial outliers in the prediction. ASSD quantifies the average distances between the volumes' surface. Precision quantifies the proportions of correctly classified cases within a class, while Recall corresponds to the proportion of positive cases correctly identified as such.

3.3. Image pre- and post-processing

Since MRI acquisitions were acquired from different centers and configurations (Winzeck et al., 2018), for each patient we re-sized all maps to a common volume of dimension of $256 \times 256 \times 32$. Afterwards, the ADC maps were clipped between $[0, 2600] \times 10^{-6} mm^2/s$ and the Tmax maps were clipped to $[0, 20s]$, since values beyond these ranges are known to be biologically meaningless (McKinley et al., 2017). Finally, a linear scaling was applied across all maps, to the range $[0, 255]$. The images are resized to its original size, after we perform the prediction.

We applied a morphological filtering as post-processing, but since the final stroke lesion presents a wide range of lesion volumes (Winzeck et al., 2018), we removed only small connected components with less than 250 voxels. This step was kept fixed for all the evaluated models.

3.3.1. Data augmentation

Data augmentation can be used to increase the number of training samples and reduce over-fitting (Krizhevsky et al., 2012). Due to the relatively small size of the training dataset, we employed

Table 1
Model training parameters for the unsupervised and supervised functional blocks.

Functional Block	Parameter	Description
Unsupervised	Optimizer	SGD with momentum ($lr = 1 \times 10^{-5}$)
	Patch shape	$7 \times 7 \times 3$
	Batch size	32
Supervised	Optimizer	ADAM ($lr = 1 \times 10^{-5}$)
	Patch shape	84×84
	Batch size	4

artificial data augmentation in the supervised portion of our proposal. For each sample, we applied rotations of 90° , 180° , 270° .

3.4. Settings and model training

Unsupervised functional block The unsupervised functional block was trained by optimizing the negative log-likelihood of the data. However, since computing the gradient is generally intractable, we performed the training by approximating the gradient with Contrastive Divergence with one step of alternating Gibbs sampling (Hinton, 2012). The training process of an RBM can be difficult if one tries to learn the parameter σ_i of the energy function, which corresponds to the standard deviation of the Gaussian noise of a visible node i (Hinton, 2012). According to Hinton (2012), we normalize each component of the data with zero mean and unit variance, and define $\sigma_i = 1$. In Table 1, we present the settings used for training the unsupervised model.

For training each RBM, we randomly extract 3D patches of shape $7 \times 7 \times 3$ from the respective input set of MRI maps, \mathcal{C} . Then, the 3D patches are reshaped into a 1D vector and fed into the visible layer of the RBM, having an input of size $m = 7 \times 7 \times 3 \times |\mathcal{C}|$, as shown in Fig. 3. After training, we extract features from the NReLU units noise-free activations. These units exhibit intensity equivariance when the bias has zero value, and they are noise free units (Nair and Hinton, 2010). Due to the large number of extracted feature maps ($|\mathcal{N}_1| = |\mathcal{N}_2| = 600$), we perform a feature selection step, as described in Section 2.2, where $\mathcal{M}_1 = \mathcal{M}_2 = 6$. The most appropriate cardinality of \mathcal{M} is discussed in Section 4.1.1.

Supervised functional block As for the supervised functional block, the complete settings of the training are given in Table 1. For each subject, 350 patches were randomly sampled. The training comprehended 36 subjects, while the remaining 7 subjects were used for validation. The settings were optimized through cross-validation in a previous work (Pinto et al., 2018a). For training, we used soft Dice loss function (Milletari et al., 2016). It is defined as:

$$\text{Soft Dice loss} = \frac{\sum_i^{|V|} p_i g_i}{\sum_i^{|V|} p_i^2 + \sum_i^{|V|} g_i^2}. \quad (2)$$

In the soft dice loss, the sum occurs over the set V of voxels belonging to the predicted output patch, where $p_i \in \mathcal{P}$ denotes the probability of a voxel i in the output patch and $g_i \in \mathcal{G}$ corresponds to the respective ground-truth label voxel.

The method was implemented using Keras with Tensorflow backend, in a workstation equipped with a GTX 1080 Ti 11 GB. Prediction time takes around 20s per patient.¹

¹ Additional details of setting and model training are provided in the supplementary material. Also, the source code for reproducing the segmentations, the models' weights and segmentations can be found at: https://github.com/apinto92/stroke_prediction.git.

4. Results and discussion

In this section, we discuss the impact of the main contributions, namely, the incorporation of unsupervised learning with supervised learning and the Gated Recurrent blocks. Then, we compare our method with the state of the art in ISLES 2017 Challenge. Finally, we delve on the difficulty of predicting the final infarct stroke lesion.

4.1. Ablation study

The ablation study aims to gradually measure the importance of the main components and consequently assert on the contribution of each component to the overall performance. Thus, we start by evaluating the importance of the unsupervised feature generator and the proposed input grouping. After, we focus on the use of the Gated Recurrent block and the choice of the dimensionality of the spatial context.

4.1.1. Unsupervised feature generation

We hypothesize that grouping the parametric MRI maps according to their physical meaning and encoding each group with an RBM has potential to extract better features to characterise the stroke lesion and the blood haemodynamics. We perform several experiments to corroborate this working hypothesis. In all experiments, the parametric MRI maps are also used as input to the supervised block. The results are presented in Table 2. Fig. 4 presents feature maps encoded by the RBMs and the respective MRI maps.

Grouping all parametric MRI maps in a single group

We considered, first, the effect of encoding all parametric maps using a single RBM. We varied the number of selected features from the RBM, observing that in all cases, the average Dice score is equal or lower than using only the parametric maps as input to the supervised block. Also, using 12 features presented the lowest average Dice score. The use of 3 or 6 obtained the same average Dice score, having the second a lower average Hausdorff distance. So, based on the metrics, we may conclude that there is no clear gain in using the features generated by the RBM, at least, when we encode all the parametric maps with a single RBM.

Since, the selection of 6 features also includes the previous top 3 features, we compared the normalized mutual information between them. As shown in Fig. 4, the top 3 features have low values of normalized mutual information in relation to the additional 3 features, which indicates that there is additional information. For this reason, we chose 6 as the number of features in the subsequent experiments.

Grouping parametric MRI maps according to the subjacent physical meaning

In this experiment, we grouped the parametric maps according to their underlining physical meaning together with ADC map in each group. Each group was encoded with an RBM. Comparing isolatedly the use of each group of features, we verify that $\text{RBM}_{\text{Lesion}}$ had a higher average Dice score compared to using only the parametric maps as input to the supervised block. The increase in the average Dice score was obtained by a higher average Recall. Also, we observe an improvement in all distance metrics. The experiment of using $\text{RBM}_{\text{Haemo}}$ presented the lowest average Dice and Recall, as well as higher average distance metrics. However, $\text{RBM}_{\text{Haemo}}$ presented higher average Precision, contrary to $\text{RBM}_{\text{Lesion}}$, which motivated the study on the combination of features from $\text{RBM}_{\text{Haemo}}$ with $\text{RBM}_{\text{Haemo}}$ besides the parametric maps. The results of this experiment are presented in Table 2. We may observe that this combination obtained the highest average Dice and Precision, as well as the lowest average distance metrics. However, this improvement could have been originated from the combination of maps according to a specific common property,

Table 2

Results obtained with different configurations of the unsupervised feature generator block in ISLES 2017 testing set. Each metric represents the mean \pm standard deviation. Underlined values correspond to the highest mean.

Unsupervised Block	Supervised Block		Dice	HD	ASSD	Precision	Recall
	FCN	G-RNN					
-	U-Net	LSTM	0.30 \pm 0.21	36.58 \pm 16.62	6.96 \pm 5.08	0.30 \pm 0.26	0.55 \pm 0.31
RBM _{single} (3 Feat.)	U-Net	LSTM	0.30 \pm 0.21	38.93 \pm 18.80	6.55 \pm 4.22	0.29 \pm 0.24	0.61 \pm 0.31
RBM _{single} (6 Feat.)	U-Net	LSTM	0.30 \pm 0.21	36.94 \pm 19.19	6.72 \pm 4.43	0.29 \pm 0.24	0.59 \pm 0.31
RBM _{single} (12 Feat.)	U-Net	LSTM	0.28 \pm 0.20	41.07 \pm 18.67	6.81 \pm 3.88	0.24 \pm 0.21	0.65 \pm 0.30
RBM _{Haemo}	U-Net	LSTM	0.28 \pm 0.24	38.50 \pm 22.78	11.09 \pm 14.79	0.35 \pm 0.30	0.44 \pm 0.34
RBM _{Lesion}	U-Net	LSTM	0.31 \pm 0.21	35.38 \pm 15.75	6.44 \pm 4.43	0.30 \pm 0.24	0.59 \pm 0.30
RBM _{Lesion} + RBM _{Haemo}	U-Net	LSTM	<u>0.38</u> \pm 0.22	<u>29.21</u> \pm 15.04	<u>5.52</u> \pm 5.06	<u>0.41</u> \pm 0.26	0.53 \pm 0.29
Two-RBM _{Mixed}	U-Net	LSTM	0.27 \pm 0.21	40.89 \pm 14.63	6.92 \pm 3.64	0.25 \pm 0.23	<u>0.68</u> \pm 0.28
Three-RBM _s ^a	U-Net	LSTM	0.35 \pm 0.23	29.32 \pm 14.33	5.27 \pm 3.54	0.34 \pm 0.27	0.59 \pm 0.30

^a RBM_{Haemo/Less} + RBM_{Lesion/Less} + RBM_{ADC}

Table 3

Results obtained when considering the Gated Recurrent block with and without the unsupervised learning block with ISLES 2017 testing set. Each metric represents the mean \pm standard deviation. Underlined values correspond to the highest mean.

Unsupervised Block	Supervised Block		Dice	HD	ASSD	Precision	Recall
	FCN	G-RNN					
-	U-Net	-	0.30 \pm 0.21	38.83 \pm 21.10	7.08 \pm 5.15	0.26 \pm 0.23	<u>0.64</u> \pm 0.30
	U-Net	LSTM	0.30 \pm 0.21	36.58 \pm 16.62	6.96 \pm 5.08	0.30 \pm 0.26	0.55 \pm 0.31
RBM _{Lesion} + RBM _{Haemo} [3D]	U-Net	-	0.32 \pm 0.23	34.09 \pm 16.51	7.60 \pm 7.14	0.35 \pm 0.27	0.48 \pm 0.32
	U-Net	LSTM	<u>0.38</u> \pm 0.22	<u>29.21</u> \pm 15.04	<u>5.52</u> \pm 5.06	<u>0.41</u> \pm 0.26	0.53 \pm 0.29

subject physical meaning of the parametric maps, in each group, or because we reduced the number of maps from 6 to 3 in each group. And this reduction could have allowed a better training of the RBM. So, we performed two complementary experiments. In the first experiment, we formed two groups with similar size, but we randomly chose the parametric maps to include in each group. In the second experiment (Three-RBMs), we changed the groups of MRI maps encoded in RBM_{Lesion} and RBM_{Haemo} by removing the ADC map from each one. These two new groups were encoded in RBM_{Lesion/Less} and RBM_{Haemo/Less}, respectively. The ADC was separately encoded in RBM_{ADC}. As presented in Table 2, the first experiment presented the lowest average Dice score and higher average distance metrics, while the second experiment attained the second highest average Dice score, thus showing the importance of splitting the parametric perfusion maps and including the ADC map in both the RBM_{Haemo} and RBM_{Lesion}.

Considering these experiments together, we may draw some conclusions. First, although CNNs are very effective in generating features from raw data, they can generate even better features if rich and complementary information is provided. A similar conclusion was inferred by Oliveira et al. (2018) that observed improvement when the coefficients of the Wavelet were added as input in the problem of retinal vessel segmentation. Here, we observe a similar effect, but using the encoding provided by an RBM trained unsupervisedly for the problem of stroke lesion prediction. Second, at least to the problem of stroke lesion prediction, when we have data with different latent factors and we are able to group it, according to those factors, then there is potential to extract complementary information from each group, but to mix them all together can be detrimental.

4.1.2. Context aggregation based on gated recurrent blocks

In medical imaging segmentation, which is similar to our problem of inferring the extension of the lesion 90 days ahead, the use of a cascade of convolutional layers to elaborate the features is the prevalent practise. However, as discussed previously, Gated-RNN layers are able to capture long distance spatial relations among input voxels, so we performed some experiments to evaluate its contribution. The results are presented in Table 3.

Analysing Table 3, we verify that when we just had parametric maps as input to the supervised block, adding a LSTM layer increased the average Precision, but the average Recall decreased, resulting in the same average Dice score. But, a different behaviour is observed when we added the features computed from the RBMs. In this scenario, we verify that using only CNN layers improved over having just parametric maps, which came by a higher average Precision. However, when we add the LSTM, we have an even higher improvement, which is observed in a larger increase in the average Precision, and a decrease in the average distance metrics.

Based on these experiments, we may conclude that the CNN layers were able to extract additional information from the RBM features; however, at least to the problem of inferring the extension of the lesion months ahead, long and local distance spatial relations among input voxels introduced by Gated RNN was critical to reduce the detection of false positives, increasing substantially the average Dice score by 6%.

4.1.3. Spatial context: 2D or 3D?

MRI images are 3D by nature, so the use of 3D filters would allow capturing more context, which has the potential to provide better prediction. Since 2D filters are confined to a plane, unnatural discontinuous contour may occur in the perpendicular axis. However, as presented previously, the resolution of MRI images in ISLES dataset is not equal in all axis, being coarser along the axial axis. So, we studied the effect of the spatial context in our architecture. As we have two blocks, unsupervised and supervised blocks, the effect on each one was evaluated separately. The results are presented in Table 4. Considering the results, we observe that using 2D patches in both blocks has lower average Dice score, than using only the parametric maps as input (baseline), because the increase in the average Precision was not enough to compensate the drop in the average Recall. Using 3D patches for both blocks had the same performance as our baseline. However, when we used 3D patches for the RBM but 2D blocks for the U-Net block, we improved over our baseline. This is the model with the highest average Dice score without LSTM. So, we may conclude that for our architecture, larger context using 3D patches was more effective for encoding features in the unsupervised block, while 2D patches

Table 4

Evaluation metrics obtained with different spatial context configurations in the unsupervised and supervised learning blocks in ISLES 2017 testing set. Each metric represents the mean \pm standard deviation. Underlined values correspond to the highest mean.

Unsupervised Block	Supervised Block		Dice	HD	ASSD	Precision	Recall
	FCN	G-RNN					
RBM _{Lesion} + RBM _{Haemo} [2D]	U-Net [2D]	-	0.27 \pm 0.23	36.35 \pm 14.89	9.14 \pm 12.35	0.31 \pm 0.28	0.53 \pm 0.34
RBM _{Lesion} + RBM _{Haemo} [3D]	U-Net [2D]	-	0.32 \pm 0.23	34.09 \pm 16.51	7.60 \pm 7.14	0.35 \pm 0.27	0.48 \pm 0.32
	U-Net [3D]	-	0.30 \pm 0.21	34.17 \pm 14.86	6.16 \pm 3.82	0.32 \pm 0.27	0.54 \pm 0.30

Table 5

Results obtained by our proposal and baseline method in the private dataset. Each metric is represented by the mean \pm standard deviation. Underlined values correspond to the highest mean, while bold values represent statistically significant values (p -value < 0.05).

Unsupervised Block	Supervised Block		Dice	HD	ASSD	Precision	Recall
	FCN	G-RNN					
-	U-Net [2D]	-	0.32 \pm 0.18	32.58 \pm 20.09	5.17 \pm 3.34	0.31 \pm 0.26	0.68 \pm 0.28
RBM _{Lesion} + RBM _{Haemo} [3D]	U-Net [2D]	LSTM	0.36 \pm 0.18	<u>26.68</u> \pm 15.60	<u>3.88</u> \pm 2.17	0.38 \pm 0.30	0.68 \pm 0.27

Table 6

Published methods in ISLES 2017 testing dataset and our proposal. Each metric is represented by the mean \pm standard deviation. Underlined values correspond to the highest mean.

		Dice	HD	ASSD	Precision	Recall	
Ensemble	Mok et al.*	0.32 \pm 0.23	40.74 \pm 27.23	8.97 \pm 9.52	0.34 \pm 0.27	0.39 \pm 0.27	
	Kwon et al.*	0.31 \pm 0.23	45.26 \pm 21.04	7.91 \pm 7.31	0.36 \pm 0.27	0.45 \pm 0.30	
	Robben et al.*	0.27 \pm 0.22	37.84 \pm 17.75	6.72 \pm 4.10	<u>0.44</u> \pm 0.32	0.39 \pm 0.31	
	Pisov et al.*	0.27 \pm 0.20	49.24 \pm 32.15	9.49 \pm 10.56	0.31 \pm 0.27	0.39 \pm 0.29	
Single Model	Monteiro et al.*	0.30 \pm 0.22	46.60 \pm 17.50	6.31 \pm 4.05	0.34 \pm 0.27	0.51 \pm 0.30	
	Pinto et al. (2018a)	0.29 \pm 0.21	41.58 \pm 22.04	7.69 \pm 5.71	0.21 \pm 0.21	<u>0.66</u> \pm 0.29	
	Lucas et al.*	0.29 \pm 0.21	33.85 \pm 16.82	6.81 \pm 7.18	0.34 \pm 0.26	0.51 \pm 0.32	
	Choi et al.*	0.28 \pm 0.22	43.89 \pm 20.70	8.88 \pm 8.19	0.36 \pm 0.31	0.41 \pm 0.31	
	Niu et al.*	0.26 \pm 0.20	48.88 \pm 11.20	6.26 \pm 3.02	0.28 \pm 0.25	0.56 \pm 0.26	
	Sedlar et al.*	0.20 \pm 0.19	58.30 \pm 20.02	11.19 \pm 9.10	0.23 \pm 0.24	0.40 \pm 0.29	
	Rivera et al.*	0.19 \pm 0.16	63.58 \pm 18.58	11.13 \pm 7.89	0.27 \pm 0.25	0.21 \pm 0.17	
	Islam et al.*	0.19 \pm 0.18	64.15 \pm 28.51	14.17 \pm 15.80	0.29 \pm 0.28	0.25 \pm 0.25	
	Chengwei et al.*	0.18 \pm 0.17	65.95 \pm 25.94	9.22 \pm 6.99	0.37 \pm 0.30	0.21 \pm 0.23	
	Yoon et al.*	0.17 \pm 0.16	45.23 \pm 19.14	12.43 \pm 11.01	0.23 \pm 0.27	0.36 \pm 0.32	
	Baseline		0.30 \pm 0.21	36.58 \pm 16.62	6.96 \pm 5.08	0.30 \pm 0.26	0.55 \pm 0.31
	Proposal		<u>0.38</u> \pm 0.22	<u>29.21</u> \pm 15.04	<u>5.52</u> \pm 5.06	0.41 \pm 0.26	0.53 \pm 0.29

*Methods presented in [Winzeck et al. \(2018\)](#), whose results were retrieved from [SMIR Online Platform \(2017\)](#).

were better suited for encoding features in the supervised U-Net-based block.

4.2. Private dataset

To further evaluate the generalization capacity of our proposal, we tested it on a private dataset and compare it with the baseline method. [Table 5](#) presents the results obtained by the two methods.

On the overall, our proposal was capable of surpassing the baseline model, attaining a higher average Dice, Precision and distance metrics, which were statistically significant (Wilcoxon Signed Ranked test with p -value < 0.05). Comparing to ISLES 2017 testing set, there was a slight decrease in performance. This could be explained by the shift on the intensity distribution of the MRI maps, due to different acquisition protocols or the differences in the preprocessing step.

4.3. State-of-the-art: ISLES 2017 challenge

The results of published methods for final infarct stroke lesion prediction using ISLES 2017 testing set ([Winzeck et al., 2018](#)), together with our baseline and proposal methods are presented in [Table 6](#). The metrics were computed by the online platform, so the ground-truth data, which was manually delineated based on a follow-up T2 MRI acquisitions, are not disclosed for public access.

Considering the results, we observe that our baseline is competitive with an average Dice, being among the top 3 methods,

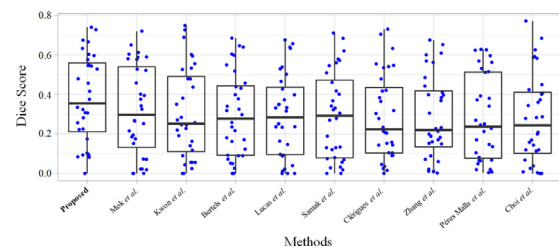


Fig. 5. Boxplot of the top-10 ranking methods ordered by average Dice score in ISLES 2017 testing set.

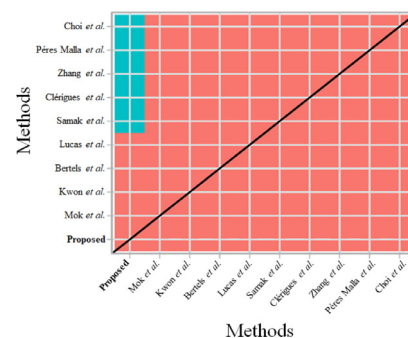


Fig. 6. One-side Wilcoxon signed rank test in ISLES 2017 testing set. Statistically significant tests are marked with the blue colour, while the red colour designates statistically non-significant tests.

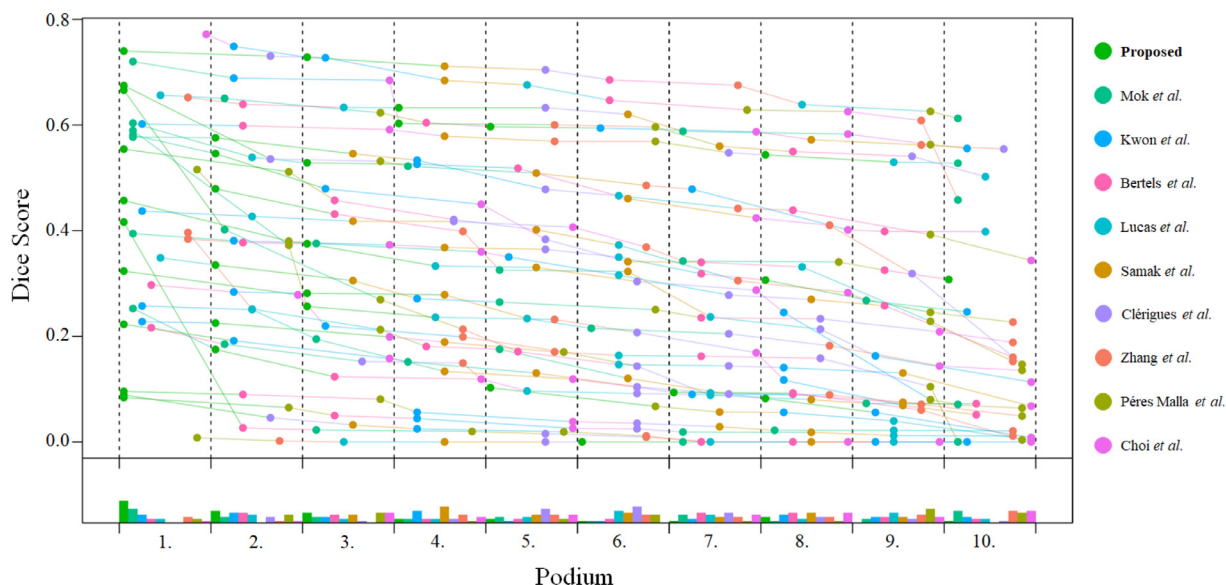


Fig. 7. Podium plot of each testing case in ISLES 2017. For each ISLES 2017 testing subject, defined by a coloured line with circles, the Podium plot orders decreasingly the Dice score obtained by each of the top 10 methods that are represented by coloured circles.

and surpassing the ensemble methods of [Pisov et al. \(2017\)](#) and [Robben and Suetens \(2017\)](#). Our proposed method presented the lowest distance metrics among all methods, especially for the Hausdorff distance. It obtained the second-best average Precision score, being surpassed by [Robben and Suetens \(2017\)](#). The authors proposed the integration of meta-data information, using a two-pathway 3D network in an ensemble; however, our experiments did not indicate any improvement using 3D patches for the U-Net, at least for our architecture. So, this improvement could have come from a combination of the effect of the ensemble and the meta-data. But we note that their method presented a lower average Recall, which explains their lower average Dice score. Regarding the average Recall score, our method was fourth, but when we consider the top 3 methods, specially [Pinto et al. \(2018a\)](#), we conclude that it was obtained with a much lower average Precision, which means that to increase the true positive detections, they had to increase substantially the false positives. So, comparing with the state of the art, our method presented a better balance between Precision and Recall, which reflected into a higher average Dice score.

Based on the results, we may conclude that the use of complementary features provided by the RBMs and the use of LSTM for a larger context allowed our baseline to surpass current state-of-the-art methods.

Results from Challenge R Benchmark

The SMIR platform of ISLES 2017 provides a weekly benchmark report of the current top-10 methods in the testing set, according to the average Dice score. So, some of the methods may not be published, lacking a description on their implementation, and, for this reason, were not included in the previous discussion.

[Fig. 5](#) presents the boxplots of each method considered in the report.

We observe that the top-10 methods failed to predict the lesion of one or more cases (lowest outliers), which may indicate the degree of complexity of predicting infarct stroke lesion 90 days ahead in ISLES 2017 Challenge dataset. But we verify that our method is the only one to have the first quartile above 0.20 in the Dice score. In [Fig. 6](#) we have the significance maps of the pairwise significant test with one-side Wilcoxon signed rank test (p -value = 0.05), showing that our method reached higher Dice score statistically

significant when compared with other five ranked methods of the top-10.

[Fig. 7](#) shows the *podium* plot of each method for each case in the testing set, and its ranking. We observe that our proposal is the method, which ranked first most of the times, as well as second and third. Also, when we consider the methods ranked below fourth, our method is in general among those with the lowest counts. Analysing the cases individually, we note two trends, for some cases all methods presented similar performance, while for others, we find a large variation from the first to the other methods. The first trend may be found in the most difficult case, where all methods had zero or a close value for the Dice score. In the second trend, we observe that our method is ranked as first most of the cases.

Based on the results of the benchmark, we may infer that our method is competitive among current state of the art, presenting the highest average Dice score and lowest average distance score. Considering the ablation study, this performance was attained due to the combination of adding extra features obtained by encoding the parametric maps with RBMs, according to the underlining physical meaning, and the elaboration provided by the long context of the LSTM layers.

5. Conclusions

In this work, we present a deep learning approach for predicting the final stroke lesion, based on unsupervised and supervised learning. We proposed to group the input maps according to the underlying physical principle behind their creation, namely, the time-resolved perfusion maps (Tmax, TTP, MTT), and the blood-flow-dynamic related maps (rCBF, rCBV). Each group was encoded using an unsupervised model to obtain structural features specific to its underlying physical principle. These structural features together with the standard parametric maps were fed to a supervised model to learn features conditioned on the label, which in our problem, means to condition on the results of the medical intervention – lesion at 90-days follow-up. We also investigated the use of Gated Recurrent Neural Networks to provide long spatial context, which were critical in relating the structural features to the information on input parametric maps. Our results showed that

either the encoding or the long spatial context improved over our baseline. Also, these two together interacted positively increasing the performance when considering separately each one.

When evaluating our proposal on ISLES 2017 testing dataset, we observe a prediction improvement over current state-of-the-art methods. The proposed method obtained the first place in Dice and also in HD and ASSD.

Recent works (Pinto et al., 2018a; Robben et al., 2020) have shown the importance of clinical meta-data to predict the final stroke lesion in different revascularization scenarios. So as future work, we aim to study how such meta-data (i.e. TICI score) could be incorporated in our architecture, to consolidate the impact of the clinical intervention and to further improve the 90-day lesion prediction.

Declaration of Competing Interest

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

CRedit authorship contribution statement

Adriano Pinto: Conceptualization, Methodology, Writing - original draft, Software. **Sérgio Pereira:** Investigation, Formal analysis, Software, Writing - review & editing. **Raphael Meier:** Data curation. **Roland Wiest:** Visualization, Investigation. **Victor Alves:** Supervision. **Mauricio Reyes:** Supervision, Validation, Writing - review & editing. **Carlos A. Silva:** Supervision, Writing - review & editing, Software, Resources.

Acknowledgments

Adriano Pinto was supported by a scholarship from the **Fundação para a Ciência e Tecnologia (FCT)**, Portugal (scholarship number [PD/BD/113968/2015](#)). This work was supported by **FCT** national funds, under the national support to R&D units grant, through the reference project [UIDB/04436/2020](#) and [UIDP/04436/2020](#).

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.media.2020.101888](#).

References

Badrinarayanan, V., Kendall, A., Cipolla, R., 2015. SegNet: a deep convolutional encoder-decoder architecture for image segmentation. arXiv:1511.00561.

Bauer, S., Gratz, P.P., Gralla, J., Reyes, M., Wiest, R., 2014. Towards automatic MRI volumetry for treatment selection in acute ischemic stroke patients. In: Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE. IEEE, pp. 1521–1524.

Berkhemer, O.A., Jansen, I.G., Beumer, D., Fransen, P.S., Van Den Berg, L.A., Yoo, A.J., Lingsma, H.F., Sprengers, M.E., Jenniskens, S.F., Lycklama à Nijeholt, G.J., et al., 2016. Collateral status on baseline computed tomographic angiography and intra-arterial treatment effect in patients with proximal anterior circulation stroke. *Stroke* 47 (3), 768–776.

Butcher, K., Emery, D., 2010. Acute stroke imaging Part I: fundamentals. *Can. J. Neurol. Sci.* 37 (1), 4–16.

Butcher, K., Emery, D., 2010. Acute stroke imaging Part II: the ischemic penumbra. *Can. J. Neurol. Sci.* 37 (1), 17–27.

Choi, Y., Kwon, Y., Lee, H., Kim, B.J., Paik, M.C., Won, J.-H., 2016. Ensemble of deep convolutional neural networks for prognosis of ischemic stroke. In: International Workshop on Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries. Springer, pp. 231–243.

Coutts, S.B., Simon, J.E., Tomanek, A.I., Barber, P.A., Chan, J., Hudon, M.E., Mitchell, J.R., Frayne, R., Eliasziw, M., Buchan, A.M., et al., 2003. Reliability of assessing percentage of diffusion-perfusion mismatch. *Stroke* 34 (7), 1681–1683.

El Tawil, S., Muir, K.W., 2017. Thrombolysis and thrombectomy for acute ischaemic stroke. *Clin. Med.* 17 (2), 161–165.

Gonzalez, R., Hirsch, J., Koroshetz, W., Lev, M., Schaefer, P., 2007. Acute ischemic stroke: imaging and intervention. *Am. J. Neuroradiol.* 28 (8), 1622.

González, R.G., Hirsch, J.A., Koroshetz, W., Lev, M.H., Schaefer, P.W., 2011. *Acute Ischemic Stroke*. Springer.

Grysiwicz, R.A., Thomas, K., Pandey, D.K., 2008. Epidemiology of ischemic and hemorrhagic stroke: incidence, prevalence, mortality, and risk factors. *Neurol. Clin.* 26 (4), 871–895.

Higashida, R.T., Furlan, A.J., Roberts, H., Tomsick, T., Connors, B., Barr, J., Dillon, W., Warach, S., Broderick, J., Tilley, B., et al., 2003. Trial design and reporting standards for intraarterial cerebral thrombolysis for acute ischemic stroke. *J. Vasc. Intervent. Radiol.* 14 (8), E1–E31.

Hinton, G.E., 2012. A practical guide to training restricted Boltzmann machines. In: *Neural Networks: Tricks of the Trade*. Springer, pp. 599–619.

Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780.

Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M., 2012. *Fsl. Neuroimage* 62 (2), 782–790.

Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B., 2017. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Med. Image Anal.* 36, 61–78.

Kemmling, A., Flottmann, F., Forkert, N.D., Minnerup, J., Heindel, W., Thomalla, G., Eckert, B., Knauth, M., Psychogios, M., Langner, S., Fiehler, J., 2015. Multivariate dynamic prediction of ischemic infarction and tissue salvage as a function of time and degree of recanalization. *J. Cereb. Blood Flow Metab.* 35 (9), 1397–1405.

Kistler, M., Bonaretti, S., Pfahrer, M., Niklaus, R., Büchler, P., 2013. The virtual skeleton database: an open access repository for biomedical research and collaboration. *J. Med. Internet Res.* 15 (11).

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 1097–1105.

Labeyrie, M.-A., Turc, G., Hess, A., Hervo, P., Mas, J.-L., Meder, J.-F., Baron, J.-C., Touzé, E., Oppenheim, C., 2012. Diffusion lesion reversal after thrombolysis: a mr correlate of early neurological improvement. *Stroke* 43 (11), 2986–2991.

Liebesskind, D.S., 2003. Collateral circulation. *Stroke* 34 (9), 2279–2284.

Lucas, C., Heinrich, M. P., 2017. 2D multi-scale res-net for stroke segmentation.

Maier, O., Menze, B.H., von der Gablentz, J., Häni, L., Heinrich, M.P., Liebrand, M., Winzeck, S., Basit, A., Bentley, P., Chen, L., et al., 2017. ISLES 2015—a public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRI. *Med. Image Anal.* 35, 250–269.

McKinley, R., Häni, L., Gralla, J., El-Koussy, M., Bauer, S., Arnold, M., Fischer, U., Jung, S., Mattmann, K., Reyes, M., et al., 2017. Fully automated stroke tissue estimation using random forest classifiers (faster). *J. Cereb. Blood Flow Metab.* 37 (8), 2728–2741.

Memezawa, H., Smith, M.-L., Siesjö, B.K., 1992. Penumbra tissues salvaged by reperfusion following middle cerebral artery occlusion in rats. *Stroke* 23 (4), 552–559.

Menon, B.K., Qazi, E., Nambiar, V., Foster, L.D., Yeatts, S.D., Liebesskind, D., Jovin, T.G., Goyal, M., Hill, M.D., Tomsick, T.A., et al., 2015. Differential effect of baseline computed tomographic angiography collaterals on clinical outcome in patients enrolled in the interventional management of stroke III trial. *Stroke* 46 (5), 1239–1244.

Milletari, F., Navab, N., Ahmadi, S.-A., 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *3D Vision (3DV), 2016 Fourth International Conference on*. IEEE, pp. 565–571.

Mok, T.C., Chung, A.C., 2017. Deep adversarial networks for stroke lesion segmentation.

Monteiro, M., Oliveira, A.L., 2017. Fully convolutional neural network for 3D stroke lesion segmentation.

Nair, V., Hinton, G.E., 2010. Rectified linear units improve restricted Boltzmann machines. In: *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 807–814.

Nielsen, A., Hansen, M.B., Tietze, A., Mouridsen, K., 2018. Prediction of tissue outcome and assessment of treatment effect in acute ischemic stroke using deep learning. *Stroke* 49 (6), 1394–1401.

Niu, Y., Gong, E., Xu, J., Pauly, J., Zaharchuk, G., 2018. Improved prediction of the final infarct from acute stroke neuroimaging using deep learning. *STROKE*, vol. 49.

Oliveira, A., Pereira, S., Silva, C.A., 2018. Retinal vessel segmentation based on fully convolutional neural networks. *Expert Syst. Appl.* 112, 229–242.

Pereira, S., Meier, R., McKinley, R., Wiest, R., Alves, V., Silva, C.A., Reyes, M., 2018. Enhancing interpretability of automatically extracted machine learning features: application to a RBM-Random Forest system on brain lesion segmentation. *Med. Image Anal.* 44, 228–244.

Pereira, S., Pinto, A., Amorim, J., Ribeiro, A., Alves, V., Silva, C.A., 2019. Adaptive feature recombination and recalibration for semantic segmentation with fully convolutional networks. *IEEE Trans. Med. Imaging* 38 (12), 2914–2925. doi:10.1109/TMI.2019.2918096.

Pinto, A., McKinley, R., Alves, V., Wiest, R., Silva, C.A., Reyes, M., et al., 2018. Stroke lesion outcome prediction based on MRI imaging combined with clinical information. *Front. Neurol.* 9, 1060.

Pinto, A., Pereira, S., Meier, R., Alves, V., Wiest, R., Silva, C.A., Reyes, M., 2018. Enhancing clinical MRI perfusion maps with data-driven maps of complementary nature for lesion outcome prediction. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pp. 107–115.

Pisov, M., Belyaev, M., Krivov, E., 2017. Neural networks ensembles for ischemic stroke lesion segmentation.

- Powers, W.J., Rabinstein, A.A., Ackerson, T., Adeoye, O.M., Bambakidis, N.C., Becker, K., Biller, J., Brown, M., Demaerschalk, B.M., Hoh, B., et al., 2018. 2018 guidelines for the early management of patients with acute ischemic stroke: a guideline for healthcare professionals from the Am. Heart Assoc./Am. Stroke Assoc.. *Stroke* 49 (3), e46–e99.
- Rekik, I., Allassonnière, S., Carpenter, T.K., Wardlaw, J.M., 2012. 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.
- Robben, D., Boers, A.M., Marquering, H.A., Langezaal, L.L., Roos, Y.B., van Oostenbrugge, R.J., van Zwam, W.H., Dippel, D.W., Majoie, C.B., van der Lugt, A., et al., 2020. Prediction of final infarct volume from native ct perfusion and treatment parameters using deep learning. *Med. Image Anal.* 59, 101589.
- Robben, D., Suetens, P., 2017. Dual-scale fully convolutional neural network for final infarct prediction. *Ischemic Stroke Lesion Segmentation-ISLES Challenge 2017, Held in Conjunction with MICCAI 2017, Date: 2017/09/10-2017/09/10, Location: Quebec City, Quebec, Canada.*
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 234–241.
- Rose, S.E., Chalk, J.B., Griffin, M.P., Janke, A.L., Chen, F., McLachan, G.J., Peel, D., Zelaya, F.O., Markus, H.S., Jones, D.K., et al., 2001. MRI based diffusion and perfusion predictive model to estimate stroke evolution. *Magn. Reson. Imaging* 19 (8), 1043–1053.
- Rumelhart, D.E., McClelland, J.L., 1986. *Parallel distributed processing: explorations in the microstructure of cognition*. Volume 1. foundations.
- Scalzo, F., Hao, Q., Alger, J.R., Hu, X., Liebeskind, D.S., 2012. Regional prediction of tissue fate in acute ischemic stroke. *Ann. Biomed. Eng.* 40 (10), 2177–2187.
- SMIR Online Platform, 2017. *Ischemic stroke lesion segmentation 2017*. <https://www.smir.ch/ISLES/Start2017>. [Accessed: 2019-11-18].
- Vinh, N.X., Epps, J., Bailey, J., 2010. Information theoretic measures for clusterings comparison: variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.* 11 (Oct), 2837–2854.
- Visin, F., Ciccone, M., Romero, A., Kastner, K., Cho, K., Bengio, Y., Matteucci, M., Courville, A., 2016. ReSeg: a recurrent neural network-based model for semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 41–48.
- Winzeck, S., Hakim, A., McKinley, R., Pinto, J.A.A.D.S., Alves, V., Silva, C., Pisov, M., Krivov, E., Belyaev, M., Monteiro, M., et al., 2018. ISLES 2016 & 2017-benchmarking ischemic stroke lesion outcome prediction based on multispectral MRI. *Front. Neurol.* 9, 679.
- World Health Organization, et al., 2014. *Global Status Report on Noncommunicable Diseases 2014*. World Health Organization.
- Zeiler, M.D., Fergus, R., 2014. Visualizing and understanding convolutional networks. In: *European Conference on Computer Vision*. Springer, pp. 818–833.
- Zivelonghi, C., Tamburin, S., 2018. Mechanical thrombectomy for acute ischemic stroke: the therapeutic window is larger but still “time is brain”. *Funct. Neurol.* 33 (1), 5.