



Automatic Cerebral Hemisphere Segmentation in Rat MRI with Ischemic Lesions via Attention-based Convolutional Neural Networks

Juan Miguel Valverde¹ · Artem Shatillo² · Riccardo De Feo¹ · Jussi Tohka¹

Accepted: 22 September 2022 / Published online: 30 September 2022
© The Author(s) 2022

Abstract

We present MedicDeepLabv3+, a convolutional neural network that is the first completely automatic method to segment cerebral hemispheres in magnetic resonance (MR) volumes of rats with ischemic lesions. MedicDeepLabv3+ improves the state-of-the-art DeepLabv3+ with an advanced decoder, incorporating spatial attention layers and additional skip connections that, as we show in our experiments, lead to more precise segmentations. MedicDeepLabv3+ requires no MR image preprocessing, such as bias-field correction or registration to a template, produces segmentations in less than a second, and its GPU memory requirements can be adjusted based on the available resources. We optimized MedicDeepLabv3+ and six other state-of-the-art convolutional neural networks (DeepLabv3+, UNet, HighRes3DNet, V-Net, VoxResNet, Demon) on a heterogeneous training set comprised by MR volumes from 11 cohorts acquired at different lesion stages. Then, we evaluated the trained models and two approaches specifically designed for rodent MRI skull stripping (RATS and RBET) on a large dataset of 655 MR rat brain volumes. In our experiments, MedicDeepLabv3+ outperformed the other methods, yielding an average Dice coefficient of 0.952 and 0.944 in the brain and contralateral hemisphere regions. Additionally, we show that despite limiting the GPU memory and the training data, our MedicDeepLabv3+ also provided satisfactory segmentations. In conclusion, our method, publicly available at <https://github.com/jmlipman/MedicDeepLabv3Plus>, yielded excellent results in multiple scenarios, demonstrating its capability to reduce human workload in rat neuroimaging studies.

Keywords Hemisphere segmentation · MRI · Convolutional Neural Networks · Rodent imaging

Introduction

Rodents are widely used in preclinical research to investigate brain diseases (Carbone, 2021). These studies often utilize in-vivo imaging technologies, such as magnetic resonance imaging (MRI), to visualize brain tissue at different time-points, which is necessary for studying disease progression. MRI permits the acquisition of brain images with different contrasts in a non-invasive manner, making MRI a particularly advantageous in-vivo imaging technology. However, these images typically need to be segmented before conducting quantitative analysis. As an example, the size

of the hemispheric brain edema relative to the volume of the contralateral hemisphere is an important biomarker for acute stroke that requires accurate hemisphere segmentation (Swanson et al., 1990; Gerriets et al., 2004).

With brain edema biomarkers in mind, our work focuses on cerebral hemisphere segmentation in MRI volumes of rat brains with ischemic lesions. Segmenting these images is particularly challenging since lesions' size, shape, location, and contrast can vary even within images from the same cohort, hampering, as we show in our experiments, traditional segmentation methods. Additionally, rodents' small size makes image acquisition sensitive to misalignments, potentially producing slices with asymmetric hemispheres and particularly affecting anisotropic data. Furthermore, although neuroanatomical segmentation tools can be used to produce hemisphere masks (e.g., Schwarz et al. (2006)), these tools only work on rodent brains without lesions, as lesions alter the appearance and location of the brain structures. These difficulties have led researchers and technicians

✉ Juan Miguel Valverde
juanmiguel.valverde@uef.fi

¹ A.I. Virtanen Institute for Molecular Sciences, University of Eastern Finland, Kuopio 70150, Finland

² Charles River Discovery Services, Kuopio 70210, Finland

to annotate rodent cerebral hemispheres manually (Freret et al., 2006; McBride et al., 2015), which is laborious and time-consuming, and motivates this work.

In recent years, convolutional neural networks (ConvNets) have been widely used to segment medical images due to their outstanding performance (Bakas et al., 2018; Bernard et al., 2018; Heller et al., 2021). ConvNets can be optimized end-to-end, require no preprocessing, such as bias-field correction and costly registration, and can produce segmentation masks in real time (De Feo et al., 2021). ConvNets can also be tailored to specific segmentation problems by incorporating domain constraints and shape priors (Kervadec et al., 2019). In particular, DeepLabv3+ architecture with its efficient computation of large image regions via dilated convolutions has demonstrated excellent results on various segmentation tasks in computer vision (Chen et al., 2018b). Xie et al. (2019) utilized DeepLabv3+ for gland instance segmentation on histology images to estimate the segmentation maps and subsequently refined such estimation with a second ConvNet. Ma et al. (2019) modified DeepLabv3+ for applying style transfer to homogenize MR images with different properties. Khan et al. (2020) showed that DeepLabv3+ outperforms other ConvNets on prostate segmentation of T2-weighted MR scans.

We present and make publicly available MedicDeepLabv3+, the first method for segmenting cerebral hemispheres in MR images of rats with ischemic lesions. MedicDeepLabv3+ improves DeepLabv3+ architecture with a new decoder with spatial attention layers (Oktay et al., 2018; Wang et al., 2019) and an increased number of skip connections that facilitate the optimization. We optimized our method on a training set comprised by 51 MR rat brain volumes from 11 cohorts acquired at multiple lesion stages, and we evaluated it on a large and challenging dataset of 655 MR rat brain volumes. Our experiments show that MedicDeepLabv3+ outperformed the baseline state-of-the-art DeepLabv3+ (Chen et al., 2018b), UNet (Ronneberger et al., 2015), HighRes3DNet (Li et al., 2017), V-Net (Milletari et al., 2016), VoxResNet (Chen et al., 2018), and, particularly for skull stripping, it also outperformed Demon (Roy et al., 2018), RATS (Oguz et al., 2014), and RBET (Wood et al., 2013). Additionally, we evaluated MedicDeepLabv3+ with very limited GPU memory and training data, and our experiments demonstrate that, despite such restrictions, MedicDeepLabv3+ yields satisfactory segmentations, showcasing its usability in multiple real-life situations and environments.

Related Work

Anatomical segmentation of rodent brain MRI with lesions

Anatomical segmentation in MR images of rodents with lesions is an under-researched area; Roy et al. (2018) and De Feo et al. (2022) are the only studies that examined this

problem. Roy et al. (2018) showed that their Inception-based (Szegedy et al., 2015) skull-stripping ConvNet named ‘Demon’ outperformed other methods on MR images of mice and humans with traumatic brain injury. De Feo et al. (2022) presented an ensemble of ConvNets named MU-Net-R for ipsi- and contralateral hippocampus segmentation on MR images of rats with traumatic brain injury. Mulder et al. (2017) developed a lesion segmentation pipeline that includes an atlas-based contralateral hemisphere segmentation step. However, these hemisphere segmentations were not compared to a ground truth, and this approach is sensitive to the lesion appearance because it relies on registration.

Anatomical segmentation of rodent brain MRI without lesions

The vast majority of anatomical segmentation methods for rodent MR brain images have been exclusively developed for brains without lesions. These methods can be classified into three categories. First, atlas-based segmentation approaches, which apply registration to one or more brain atlases (Pagani et al., 2016) and, afterwards, label candidates are refined or combined with, for instance, Markov random fields (Ma et al., 2014). As these approaches heavily rely on registration, they underperform in the presence of anatomical deformations. Second, methods that group nearby voxels with similar properties. These approaches typically start by proposing one or several candidate regions, and later adjust such regions with an energy function and, optionally, shape priors. Examples of these methods include surface deformation models (Wood et al., 2013), graph-based segmentation algorithms (Oguz et al., 2014), and a more recent approach that combines blobs into a single region (Liu et al., 2020). These approaches can handle different MRI contrasts and require no registration. However, they also rely on local features, such as nearby image gradients and intensities. Thus, these methods can be very sensitive to intensity inhomogeneities, and small brain deformities. Third, machine learning algorithms that classify brain features. These features can be handcrafted, such as in (Bae et al., 2009; Wu et al., 2012) where authors employed support vector machines to classify voxels into different neuroanatomical regions based on their intensity, location, neighbor labels, and probability maps. On the contrary, deep neural networks, a subclass of machine learning algorithms, can automatically find relevant features and learn meaningful non-linear relationships between such features. Methods based on neural networks, such as pulse-coupled neural networks (Chou et al., 2011; Murugavel and Sullivan, 2009) and ConvNets (Roy et al., 2018; Hsu et al., 2020; De Feo et al., 2021), have been used in the context of rodent MRI segmentation.

Lesion segmentation of rodent brain MRI The high contrast between lesion and non-lesion voxels in certain rodent brain MR images motivated the development of

thresholding-based methods (Wang et al., 2007; Choi et al., 2018). However, these methods are not fully automatic, and they cannot be used in MR images with other contrasts, or lesions with different appearances. Mulder et al. (2017) introduced a fully-automated pipeline to segment lesions via level sets (Dervieux & Thomasset, 1980; Osher & Sethian, 1988). Images were first registered to a template, then skull stripped, and their ventricles were segmented prior to the final lesion segmentation step. Arnaud et al. (2018) framed lesion segmentation as an anomaly-detection problem and developed a pipeline that detects voxels with unusual intensity values with respect to healthy rodent brains. Valverde et al. (2020) developed the first single-step method to segment rodent brain MRI lesions using ConvNets.

Materials and Methods

MRI Data

The image data, provided by Charles River Laboratories Discovery site (Kuopio, Finland)¹, consisted of 723 MR T2-weighted brain scans of 481 adult male Wistar rats weighting between 250–300 g derived from 11 different cohorts. Rats were induced focal cerebral ischemia by middle cerebral artery occlusion for 120 minutes in the right hemisphere of the brain (Koizumi et al., 1986). MR data was acquired at multiple time-points after the occlusion; for each of the 11 cohorts, time-points were different (see Fig. 1A for details). In total, our dataset contained MR images from nine lesion stages: shams, 2h, 24h, D3, D7, D14, D21, D28, and D35. Figure 1B shows representative images of these lesion stages in approximately the same brain area. All animal experiments were conducted according to the National Institute of Health (NIH) guidelines for the care and use of laboratory animals, and approved by the National Animal Experiment Board, Finland. Multi-slice multi-echo sequence was used with the following parameters; TR = 2.5 s, 12 echo times (10–120 ms in 10 ms steps) and 4 averages in a horizontal 7T magnet. T2-weighted images were calculated as the sum of the all echoes. Eighteen coronal slices of 1 mm thickness were acquired using a field-of-view of 30x30 mm² producing 256x256 imaging matrices of resolution 117 × 117 μm. Afterwards, these coronal slices were combined into a 3D volume.

Data Preparation

The T2-weighted MRI volumes were not preprocessed (i.e., no registration, bias-field or artifact correction), and their

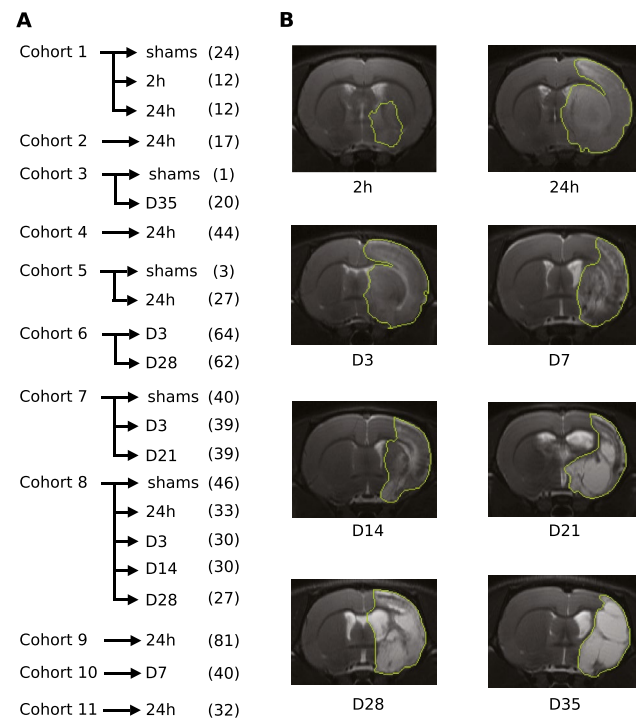


Fig. 1 **A:** Cohorts, acquisition time-points, and number of images. **B:** Example slice from each lesion stage in approximately the same brain area. The lesion border of the manual segmentation is outlined in green

intensity values were standardized to have zero mean and unit variance. Brain and contralateral hemisphere masks were annotated by several trained technicians employed by Charles River according to a standard operating procedure. These annotations did not include the cerebellum and the olfactory bulb. Finally, we computed the ipsilateral hemisphere mask by subtracting the contralateral hemisphere from the brain mask, yielding non-overlapping regions (i.e., the background, ipsilateral and contralateral hemispheres) for optimizing the ConvNets.

Train, Validation and Test Sets

We divided the MR images into a training set of 51 volumes, validation set of 17 volumes, and test set of 655 volumes. Specifically, we grouped the MR images by their cohort and acquisition time-point (Fig. 1A). From the resulting 17 subgroups, our training and validation sets comprised 3 and 1 MR images, respectively, per subgroup. Images from sham-operated animals were not included to the training and validation sets since our work focused on rat brains with lesions. The remaining 655 MR images, including shams, formed the independent test set. This splitting strategy aimed to create a diverse training set, as brain lesions have notably different T2-weighted MRI intensities depending on the lesion stage,

¹ <https://www.criver.com/products-services/discovery-services>

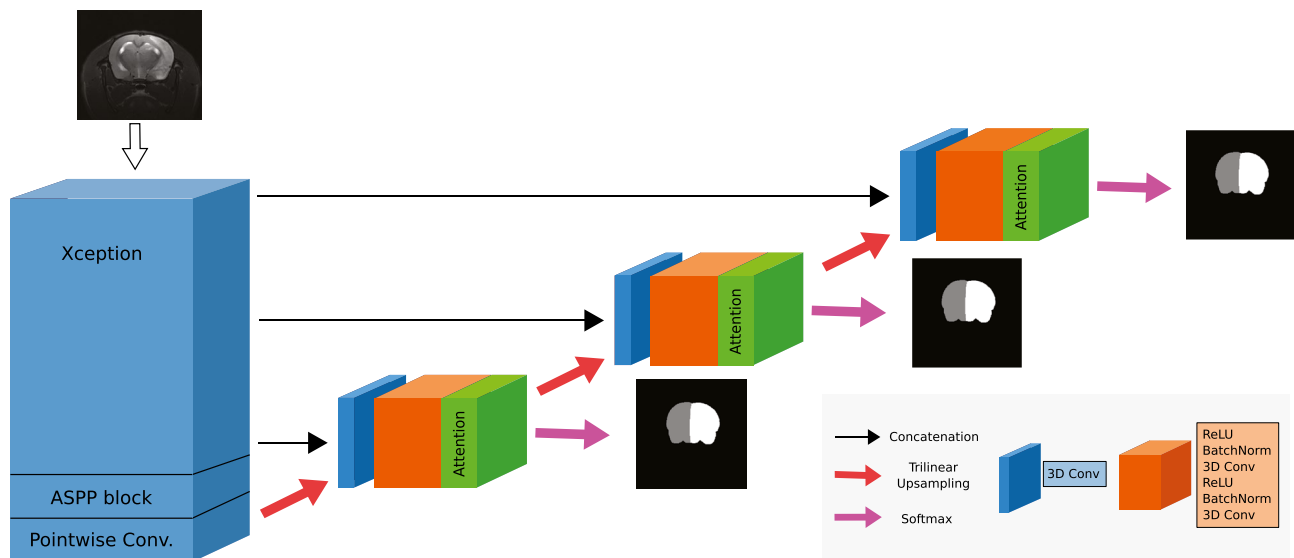


Fig. 2 MedicDeepLabv3+ architecture

and annotations can differ slightly across cohorts due to the task subjectivity and the consequent low inter-rater agreement (Mulder et al., 2017; Valverde et al., 2020).

MedicDeepLabv3+

MedicDeepLabv3+ is a 3D fully convolutional neural network (ConvNet) based on DeepLabv3+ (Chen et al., 2018b) and UNet (Ronneberger et al., 2015). We chose DeepLabv3+ because of its excellent performance in semantic segmentation tasks, and we modified its last layers to resemble more closely to UNet, which is an architecture widely used in medical image segmentation. DeepLabv3+ first employs Xception (Chollet, 2017) to transform the input and reduce its dimensionality, and then it upsamples the transformed data, twice, by a factor of four. MedicDeepLabv3+ replaces these last layers (i.e., the decoder) with three stages of skip connections and convolutional layers, and, as we describe below, it incorporates custom spatial attention layers, enabling deep supervision (see Fig. 2).

Encoder

MedicDeepLabv3+ stacks several $3 \times 3 \times 3$ convolutional layers, normalizes the data with Batch Normalization (Ioffe and Szegedy, 2015), and incorporates residual connections (He et al., 2016). Both batch normalization and residual connections are well established architectural components that have been shown to facilitate the optimization in deep ConvNets (Drozdzal et al., 2016; Li et al., 2018). The first layers of MedicDeepLabv3+ correspond to Xception (Chollet, 2017), which uses depthwise-separable convolutions instead

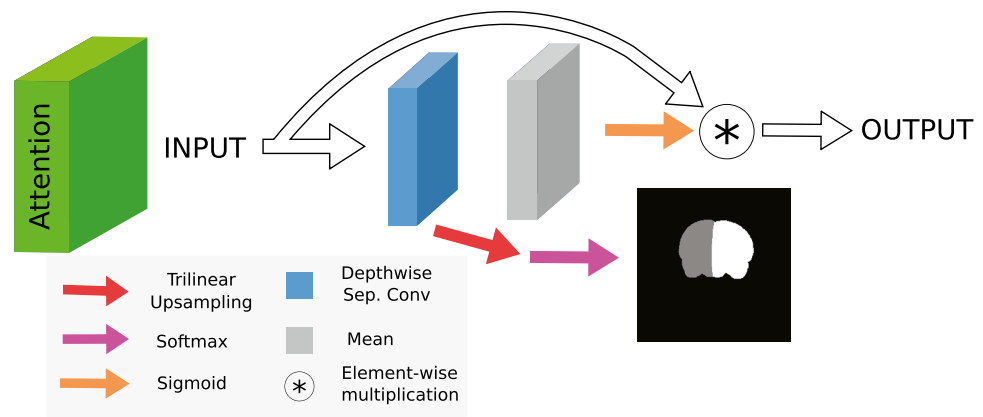
of regular convolutions. These depthwise-separable convolutions are advantageous over regular convolutions as they can decouple channel and spatial information. This is achieved by separating the operations of a regular convolution into a spatial feature learning and a channel combination step, increasing the efficiency and performance of the model (Chollet, 2017).

MedicDeepLabv3+ utilizes dilated convolutions in the last layer of Xception. Dilated convolutions sample padded input patches and multiply the non-padded values with the convolution kernel, thus, expanding the receptive field of the network (Chen et al., 2014). In other words, dilated convolutions permit to adjust the area that influences the classification of each voxel, and, as increasing this area has shown to improve model performance, we opted to employ dilated convolutions as in DeepLabv3+ (Chen et al., 2017). After the Xception backbone, DeepLabv3+'s Atrous Spatial Pyramid Pooling (ASPP) module concatenates parallel branches of dilated convolutional layers with different dilation rates and an average pooling followed by trilinear interpolation. Then, a pointwise convolution combines and reduces the number of channels. To this step, the described architecture reduces the data dimensionality by a factor of 16.

Decoder

We developed a new decoder for MedicDeepLabv3+ with more stages of skip connections and convolution blocks than DeepLabv3+. In each stage, feature maps are upsampled via trilinear interpolation and concatenated to previous feature maps from the encoder. Subsequently, $3 \times 3 \times 3$ convolutions halve the number of channels (Fig. 2, blue

Fig. 3 Spatial attention block (details in Sect. 2.4.2)



blocks), and a ResNet block (He et al., 2016) further transforms the data (Fig. 2, orange blocks). The consequent increase of skip-connections facilitates MedicDeepLabv3+ optimization (Drozdzal et al., 2016; Li et al., 2018). Importantly, DeepLabv3+ produces segmentations at $\times 4$ less resolution than the original images that, to match their size, are upsampled via interpolation. In contrast, our MedicDeepLabv3+ incorporates more convolutional layers at the end of its architecture to perform a final data transformation at the same resolution as the input.

Another key difference with respect to DeepLabv3+ is that MedicDeepLabv3+ utilizes spatial attention layers (Oktay et al., 2018; Wang et al., 2019). Attention layers behave as dynamic activation functions that first learn and then apply voxel-wise importance maps, transforming feature maps differently based on their values. Consequently, attention layers permit to learn more complex relations in the data. Furthermore, attention layers have been shown to aid the network to identify and focus on the most important features, improving performance (Fu et al., 2019). In our implementation, these layers (Fig. 3) transform the inputs with a depthwise-separable convolution and, subsequently, average the resulting feature maps. Afterwards, a sigmoid activation function transforms the data non-linearly, producing spatial attention maps with values in the range $[0, 1]$. Then, these attention maps multiply the input feature maps voxel-wise. To encourage spatial attention maps that lead to the ground truth and to further facilitate the optimization, we added a branch in the first two attention layers for generating downsampled probability maps of the segmentation masks, enabling deep supervision (Fig. 3, red and pink arrows).

Loss Function

We trained MedicDeepLabv3+ with deep supervision (Lee et al., 2015), i.e., we minimized the sum of cross entropy and Dice loss of all outputs of MedicDeepLabv3+ (Fig 2,

pink arrow). Formally, we minimized $L = \sum_{s \in S} L_{CE}^s + L_{Dice}^s$ with $S = \{1, 2, 3\}$ indicating each MedicDeepLabv3+ output (see Fig. 2). Cross entropy treats the model predictions and the ground truth as distributions

$$L_{CE} = -\frac{1}{NC} \sum_{i=1}^N \sum_{c=1}^C p_{i,c} \log(q_{i,c}), \tag{1}$$

where $p_{i,c} \in \{0, 1\}$ represents whether voxel i belongs to class c , and $q_{i,c} \in [0, 1]$ its predicted Softmax probability. $C = 3$ for background, ipsilateral and contralateral hemisphere classes, and N is the total number of voxels. Dice loss estimates the Dice coefficient between the predictions and the ground truth:

$$L_{Dice} = 1 - \frac{2}{C} \sum_{c=1}^C \frac{\sum_i p_{i,c} q_{i,c}}{\sum_i p_{i,c}^2 + q_{i,c}^2}. \tag{2}$$

Minimizing cross entropy and Dice loss is a common practice in medical image segmentation (Myronenko & Hatamizadeh, 2019; Isensee et al., 2021). Cross entropy optimization reduces the difference between the ground truth and prediction distributions. Dice loss optimization increases the Dice coefficient that we ultimately aim to maximize and it is particularly beneficial in class-imbalanced datasets (Milletari et al., 2016). Additionally, their optimization at different stages via deep supervision is equivalent to adding shortcut connections to propagate the gradients to various layers, facilitating the optimization of those layers.

Experimental Design

Metrics

We assessed the automatic segmentations with Dice coefficient (Dice, 1945), Hausdorff distance (Rote, 1991), precision, and recall. Dice coefficient measures the overlapping volume between the ground truth and the prediction

$$Dice(A, B) = \frac{2|A \cap B|}{|A| + |B|}, \quad (3)$$

where A and B are the segmentation masks. Hausdorff distance (HD) is a quality metric that calculates the distance to the misclassification located the farthest from the boundary masks. Formally:

$$d(A, B) = \max \left\{ \max_{a \in \partial A} \min_{b \in \partial B} |b - a|, \max_{b \in \partial B} \min_{a \in \partial A} |a - b| \right\}, \quad (4)$$

where ∂A and ∂B are the boundary voxels of A and B , respectively. In other words, HD provides the distance to the largest segmentation error. We provided HD values in mm and we accounted for voxel anisotropy. Finally, precision is the percentage of voxels accurately classified as brain/hemisphere, and recall is the percentage of brain/hemisphere voxels that were correctly identified:

$$Prec = \frac{TP}{TP + FP} \quad Recall = \frac{TP}{TP + FN}. \quad (5)$$

Benchmarked Methods

We compared our MedicDeepLabv3+ with DeepLabv3+ baseline (Chen et al., 2018b), UNet (Ronneberger et al., 2015), HighRes3DNet (Li et al., 2017), V-Net (Milletari et al., 2016), VoxResNet (Chen et al., 2018), Demon (Roy et al., 2018), RATS (Oguz et al., 2014), and RBET (Wood et al., 2013). Since Demon, RATS, and RBET were exclusively designed for rodent skull stripping, we computed contralateral hemisphere masks only with MedicDeepLabv3+, DeepLabv3+, UNet, HighRes3DNet, V-Net, and VoxResNet. MedicDeepLabv3+ and all the other ConvNets were optimized on our training set with Adam (Kingma & Ba, 2014) ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$), starting with a learning rate of 10^{-5} . ConvNets' hyper-parameters (e.g., activation functions, number of filters, number of convolutional layers, normalization layers, type of downsampling/upsampling layers, dilation rate in atrous convolution layers, kernel size) were set as instructed in the articles proposing the method (i.e., DeepLabv3+ (Chen et al., 2018b), UNet (Ronneberger et al., 2015), HighRes3DNet (Li et al., 2017), V-Net (Milletari et al., 2016), VoxResNet (Chen et al., 2018a), Demon (Roy et al., 2018)). As such, these hyper-parameters can be considered to be close to optimal. MedicDeepLabv3+, DeepLabv3+, HighRes3DNet, V-Net, VoxResNet and UNet were trained for 300 epochs, and Demon was trained for an equivalent amount of time. All ConvNets were either 2D or 3D; we disregarded 2.5D ConvNets (e.g., Kushibar et al. (2018)) that concatenate three orthogonal 2D images to classify the voxel at their intersection since they are more inefficient at inference time than 2D and 3D ConvNets. We ensembled three models (Dietterich, 2000) since this strategy markedly improved segmentation performance in

our previous work (Valverde et al., 2020). More specifically, we trained each ConvNet three times, separately, starting from different random initializations. Then, we formed the final segmentations based on the majority vote from the binarized outputs of the three trained models.

We conducted a grid-search for best hyperparameters for RATS and RBET. We performed the grid-search using merged training and validation sets as RATS and RBET do not involve supervised learning, thus making it possible to use also the training set for hyper-parameter tuning. Subsequently, we utilized the best-performing hyper-parameters on the test set. With RATS, computing the brain mask \hat{y} of image x requires setting three hyper-parameters: intensity threshold t , α , and rodent brain volume s , i.e., $\hat{y} = RATS(x, t, \alpha, s)$. As rat brain volumes are highly similar in adult rats, we left this hyper-parameter with its default value, $s = 1650 \text{ mm}^3$. Thus, we only optimized for the threshold t and α hyper-parameters. Since RATS assumes that all intensity values are positive integers, we employed unnormalized images with RATS. We optimized RATS hyper-parameters by maximizing the Dice coefficients in the training and validation sets:

$$\arg \max_{i, \alpha} \sum_{x \in X_{train+val}} Dice(y, RATS(x, t, \alpha, 1650)) : t = P_{\%i}, \quad (6)$$

where $Dice$ is the Dice coefficient (Eq. (3)) between the ground-truth brain mask y and RATS' output, $\alpha = 0, 1, \dots, 10$ balances the importance between gradients and intensity values, and $P_{\%i}$ is the i th percentile of x with $i = 0.01, 0.02, \dots, 0.99$. Since finding t is potentially suboptimal due to the distribution variability across images, we optimized for the i th percentile, yielding image-specific thresholds. In total, our hyper-parameter grid search in RATS comprised 1089 different parameter value combinations. For RBET, we optimized the Dice coefficient to find the optimal ellipse axes ratio $w:h:d$ with w, h, d from 0.1 to 1 in steps of 0.05, accounting for 19^3 different configurations. Note that, despite optimizing over a large number of hyper-parameter choices may increase the risk to overfit, our train, validation and test sets were derived so that $X_{train+val}$ is a good representation of X_{test} (see Sec. 2.3).

Unlike ConvNets that can be optimized to segment specific brain regions, RATS and RBET perform skull stripping, segmenting also the cerebellum and olfactory bulb that were not annotated. As these brain areas were not part of our ground truth, RATS and RBET segmentations would be unnecessarily penalized in those areas. Thus, before computing the metrics, we discarded the slices containing cerebellum and olfactory bulb. This evaluation strategy ignores potential misclassifications in the excluded slices, slightly favoring RATS and RBET.

We tested whether the difference in Dice coefficient and HD between our MedicDeepLabv3+ and the other methods was significant. The hypothesis was tested with a paired

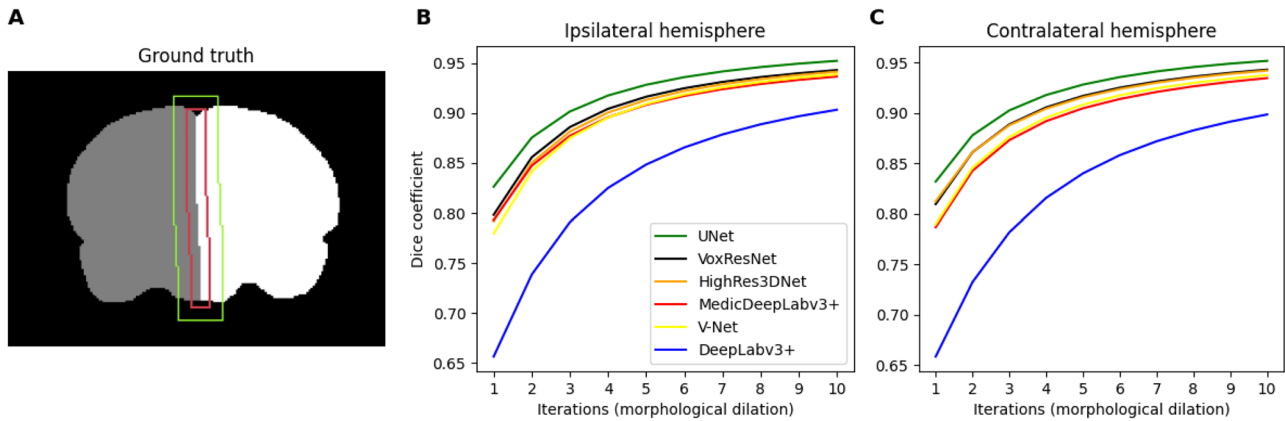


Fig. 4 A: Example of ground truth and its brain midline area after four (red) and ten (green) iterations of morphological dilation. B-C: Dice coefficients for the ipsi- and contralateral hemisphere classes in

the brain midline area with different morphological dilation iterations (brain midline area sizes)

two-sample permutation test using the mean-absolute difference as the test statistic. We considered p -values smaller than 0.05 as statistically significant.

Brain Midline Evaluation

We calculated the average Dice coefficients of contra- and ipsilateral hemispheres around the brain midline—boundary between both hemispheres (see Fig. 4A). Specifically, we considered the volume after expanding brain midline voxels in the coronal plane via morphological dilation n times, with $n = 1, 2, \dots, 10$. In contrast to brain vs. non-brain tissue boundaries, the brain midline volume is more ambiguous to annotate due to the lower intensity contrast between hemispheres, hence the importance to assess the performance in this area. This evaluation aims to supplement computing Dice coefficient and HD on the whole 3D volumes. Since most of the voxels lie within the hemisphere borders, Dice coefficients tend to be very high, and since HD might indicate the distance to a misclassification that can be easily corrected via postprocessing (e.g., a misclassification outside the brain), HD alone does not suffice to assess specific areas. Note that, similarly to RATS and RBET evaluation, this experiment computed the Dice coefficient only on the slices that were manually annotated, as finding the brain midline requires these manual annotations. Consequently, the evaluated 3D masks excluded non-annotated slices that could have false positives.

Biomarkers Based on Hemisphere Segmentation

Since the ratio between contra- and ipsilateral hemispheres volume is an important biomarker for acute stroke (Swanson et al., 1990; Gerriets et al., 2004), we compared the hemisphere

volume ratio of the ground truth with the hemisphere volume ratio of the automatic segmentations. For this, we computed the effect size via Cohen's d (Lakens, 2013) and the bias-corrected and accelerated (BCa) bootstrap confidence intervals (Efron, 1987) with 100000 bootstrap resamples. An effect size close to zero with a narrow confidence interval indicates a high similarity between automated and manual segmentation based biomarkers.

Performance with Limited GPU Memory and Data

Motivated by potential GPU memory limitations, we studied the performance and computational requirements of multiple versions of MedicDeepLabv3+ with lower capacity and, consequently, lower GPU memory usage. To investigate this, we varied the number of kernel filters in all convolutions of MedicDeepLabv3+ that determines the number of parameters. For instance, decreasing the number of kernel filters by half in the encoder also decreases the number of kernel filters in the decoder to half.

Separately, we evaluated the proposed MedicDeepLabv3+ on each cohort and time-point independently, simulating the typical scenario in rodent studies with extremely scarce annotated data. For each of the 17 groups containing no sham animals (Fig. 1A), we trained an ensemble of three MedicDeepLabv3+ on *only three images*, employed another image for validation during the optimization, and we evaluated this ensemble on the remaining holdout images from the same group.

Implementation

MedicDeepLabv3+, DeepLabv3+, UNet, HighRes3DNet, V-Net, VoxResNet, and Demon were implemented in Pytorch (Paszke et al., 2019) and were run on Ubuntu 16.04 with an Intel Xeon W-2125 CPU @ 4.00GHz processor,

Table 1 Dice coefficients, Hausdorff distances (HD), precision, and recall of the brain and contralateral hemisphere (CH) masks derived from the evaluated methods (mean \pm std). Bold: best scores. Scores that were significantly different from our MedicDeepLabv3+ were marked with * (p -value < 0.05)

	Approach	Dice	HD	Prec	Recall
Brain	MedicDeepLabv3+	0.952 \pm 0.04	1.856 \pm 0.91	0.94 \pm 0.07	0.97 \pm 0.03
	VoxResNet	0.951 \pm 0.04	2.042 \pm 1.02*	0.94 \pm 0.07	0.97 \pm 0.02
	HighRes3DNet	0.949 \pm 0.04*	1.858 \pm 1.04	0.93 \pm 0.07	0.97 \pm 0.02
	V-Net	0.948 \pm 0.04*	1.920 \pm 1.05*	0.94 \pm 0.07	0.97 \pm 0.02
	UNet	0.947 \pm 0.05*	3.477 \pm 1.20*	0.93 \pm 0.07	0.97 \pm 0.02
	DeepLabv3+	0.936 \pm 0.04*	2.149 \pm 1.02*	0.93 \pm 0.07	0.95 \pm 0.03
	Demon	0.934 \pm 0.04*	3.621 \pm 1.17*	0.92 \pm 0.07	0.96 \pm 0.02
	RATS	0.913 \pm 0.01*	2.221 \pm 0.51*	0.91 \pm 0.03	0.92 \pm 0.02
	RBET	0.781 \pm 0.10*	3.628 \pm 0.46*	0.89 \pm 0.05	0.70 \pm 0.10
CH	MedicDeepLabv3+	0.944 \pm 0.04	2.064 \pm 1.85	0.94 \pm 0.08	0.96 \pm 0.03
	VoxResNet	0.944 \pm 0.04	2.265 \pm 1.86*	0.93 \pm 0.07	0.96 \pm 0.02
	HighRes3DNet	0.942 \pm 0.04*	2.205 \pm 1.86*	0.93 \pm 0.07	0.96 \pm 0.03
	V-Net	0.940 \pm 0.04*	2.218 \pm 1.86*	0.93 \pm 0.07	0.96 \pm 0.03
	UNet	0.941 \pm 0.05*	3.689 \pm 1.64*	0.92 \pm 0.07	0.97 \pm 0.02
	DeepLabv3+	0.921 \pm 0.04*	2.411 \pm 1.80*	0.91 \pm 0.07	0.94 \pm 0.03

64 GB of memory and an NVidia GeForce GTX 1080 Ti with 11 GB of memory. MedicDeepLabv3+ and the scripts for segmenting rat MR images and to optimize new models are publicly available at <https://github.com/jmlipman/MedicDeepLabv3Plus>. These scripts are ready for use via command line interface with a single command, and users can easily adjust the number of initial filters that controls the model size, capacity, and GPU memory requirements. Additionally, we provide the optimized parameters (i.e., the weights) of MedicDeepLabv3+ at <https://github.com/jmlipman/MedicDeepLabv3Plus>.

Results

Segmentation Metrics Comparison

Our MedicDeepLabv3+ produced brain and hemisphere masks with the highest Dice coefficients (0.952 and 0.944) and precision (0.94 and 0.94), and the lowest HD (1.856 and 2.064) (see Table 1). MedicDeepLabv3+ also achieved the highest Dice coefficients in the brain and contralateral hemisphere most frequently, in 38% and 36% of the test images, respectively, followed by UNet (24% and 26%), VNet (13% and 13%), VoxResNet (13% and 12%), HighRes3DNet (11% and 12%) and the others (1% or less). In the majority of cases, MedicDeepLabv3+ produced segmentations with Dice and HD significantly better than the compared methods (see Table 1). All ConvNets performed better than RATS and RBET and, particularly, 3D ConvNets (MedicDeepLabv3+, DeepLabv3+, HighRes3DNet, V-Net, and VoxResNet) consistently yielded lower HD than 2D ConvNets (UNet, Demon). Our MedicDeepLabv3+ produced finer segmentations that were more

similar to the ground truth than the baseline DeepLabv3+ which generated masks with imprecise borders. UNet also produced segmentations with higher Dice and recall than DeepLabv3+, although UNet HD was considerably lower. HighRes3DNet, V-Net, and VoxResNet yielded slightly worse Dice coefficients and HDs than MedicDeepLabv3+. Figure 5 illustrates these results on the MR image with the highest hemispheric volume imbalance. Figure 5 shows that RBET was incapable of finding the brain boundaries; RATS produced segmentations with several holes and non-smooth borders; 2D ConvNets misclassified the olfactory bulb and cerebellum; and, in agreement with Table 1, MedicDeepLabv3+ produced the segmentation mask most similar to the ground truth. We included 17 images (one per cohort and lesion time-point) in Online Resource 1 that also corroborate the higher performance of MedicDeepLabv3+. The computation time to optimize these methods also varied notably: on average, ConvNets required 16 hours, and RATS and RBET needed six days. Furthermore, MedicDeepLabv3+ segmented the images in real time, requiring approximately 0.4 seconds per image.

Brain Midline Experiment

Regarding the brain midline area experiment (Sect. 2.5.3, Fig. 4B,C), MedicDeepLabv3+ outperformed the baseline DeepLabv3+ across different area sizes (average Dice coefficient difference of 0.07). VoxResNet, HighRes3DNet, V-Net, and MedicDeepLabv3+ yielded very similar Dice coefficients, and UNet produced the highest Dice coefficients by a small margin (average difference between UNet and MedicDeepLabv3+ of only 0.02). Additionally, Dice coefficients were similar across hemispheres regardless of the segmentation method.

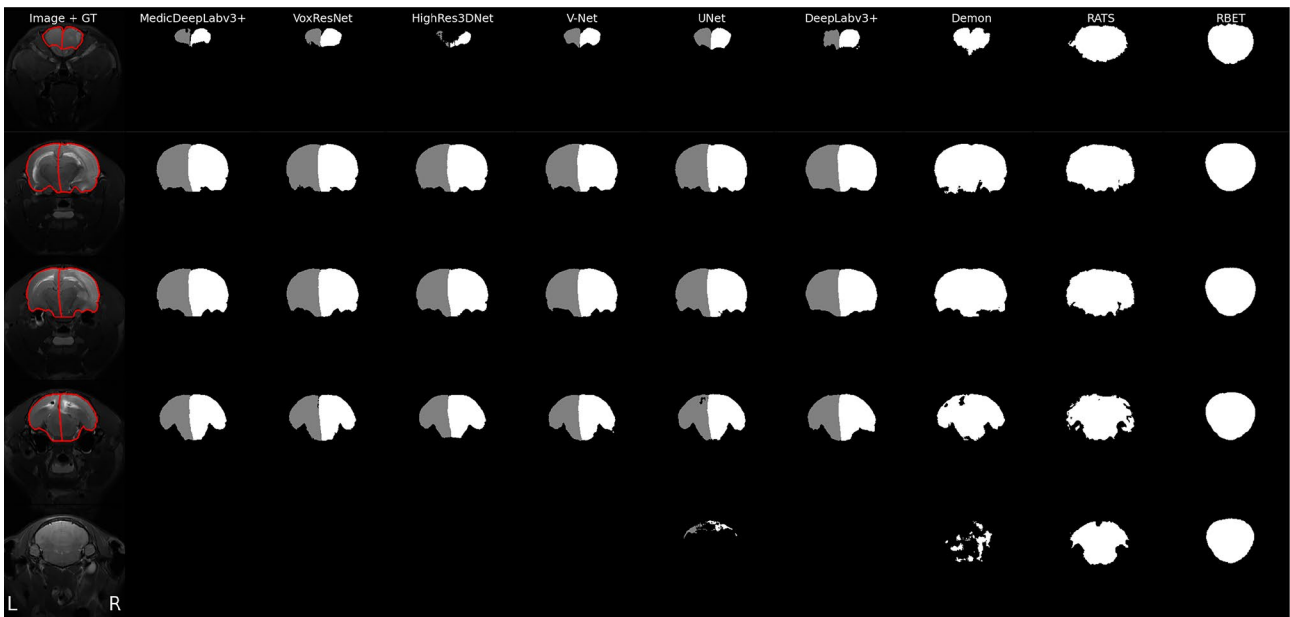


Fig. 5 T2-weighted image, ground truth and automatic segmentations of the rat with the most imbalanced hemispheric volumes

Hemispheric Ratio Experiment

The computed Cohen’s d shows that, in terms of magnitude, all methods produced hemispheric ratio distributions not too different from the ground truth (Table 2). Among these methods, MedicDeepLabv3+ and V-Net provided the smallest effect size and the most zero-centered confidence interval, with MedicDeepLabv3+’s confidence interval being narrower than V-Net’s. DeepLabv3+’s confidence interval was the largest and contained zero whereas UNet’s confidence interval was the narrowest—slightly narrower than MedicDeepLabv3+’s—and did not contain zero.

Limited Resources

Table 3 lists the characteristics, computational requirements, and performance of different versions of MedicDeepLabv3+ on the contralateral hemisphere segmentation (performance on the brain can be found in Online Resource 2). Reducing the number of parameters by decreasing the number of initial filters reduced notably the required GPU memory and training time while it barely affected MedicDeepLabv3+’s performance. For instance, reducing the number of parameters by 93.5% (from 79.1M to 5.1M) decreased the required GPU memory and training time by 72% while it decreased the Dice coefficient in the contralateral hemisphere by only 1%.

Tables 4 and 5 show the performance of MedicDeepLabv3+ optimized and evaluated on each cohort and acquisition time-point separately. In other words, for each cohort and acquisition time-point, the training set was comprised by only three images and test set size (Tables 4 and 5,

“Volumes” column) varied across the 17 groups. MedicDeepLabv3+, on average, performed slightly worse than in our first experiment that utilized 17 times more annotated data. Performance measures across these groups varied notably: in the contralateral hemisphere segmentations (Table 5) Dice coefficients ranged from 0.876 to 0.951, HD from 1.200 to 3.745, precision from 0.871 to 0.962, and recall from 0.859 to 0.967. Additionally, in agreement with our previous experiment, performance on the contralateral hemisphere was slightly lower than on the brain.

Discussion

We presented MedicDeepLabv3+, the first method for hemisphere segmentation in rat MR images with ischemic lesions. We compared MedicDeepLabv3+ performance with state-of-the-art DeepLabv3+, UNet, HighRes3DNet, V-Net, VoxResNet, and three brain extraction algorithms (Demon,

Table 2 Cohen’s d that measured the effect size and its confidence intervals

Approach	Cohen’s d	Confidence Interval
MedicDeepLabv3+	0.008	[-0.013, 0.035]
VoxResNet	-0.042	[-0.060, -0.025]
HighRes3DNet	-0.102	[-0.125, -0.080]
V-Net	0.003	[-0.042, 0.022]
UNet	-0.038	[-0.054, -0.021]
DeepLabv3+	0.050	[-0.008, 0.099]

Table 3 Comparison between multiple versions of MedicDeepLabv3+ with different capacity. Columns: proportion of kernel filters with respect to the default configuration, trainable ConvNet parameters (in millions), optimization time for 300 epochs in our workstation (see Sect. 2.5.6 for

details) in hours, maximum GPU memory required during training and evaluation, Dice and HD in the contralateral hemisphere (mean \pm std). Bold: default configuration, highest performance

Rate	Parameters	Time (h)	Mem. (train)	Mem. (eval)	Dice	HD
1	79.1M	16.2	8857 MiB	2935 MiB	0.944 \pm 0.04	2.064 \pm 1.85
0.875	60.7M	14.4	7571 MiB	2617 MiB	0.941 \pm 0.04	2.103 \pm 1.86
0.750	44.7M	12.1	6545 MiB	2319 MiB	0.941 \pm 0.04	2.118 \pm 1.86
0.625	31.1M	10.3	5619 MiB	2007 MiB	0.941 \pm 0.04	2.099 \pm 1.85
0.500	20.0M	7.8	4577 MiB	1717 MiB	0.939 \pm 0.05	2.099 \pm 1.83
0.375	11.3M	6.2	3531 MiB	1421 MiB	0.937 \pm 0.05	2.138 \pm 1.85
0.250	5.1M	4.5	2503 MiB	1121 MiB	0.933 \pm 0.05	2.078 \pm 1.82

RATS, and RBET) combining several preclinical neuroimaging studies to a large dataset of 723 rat MR volumes.

ConvNets performed markedly better and their training time was about 10 times shorter than RATS (Oguz et al., 2014) and RBET (Wood et al., 2013). The superior performance of ConvNets was not surprising, as RATS and RBET were not designed to segment brains with widely varying intensity values, such as those found in brains with lesions. This out-performance of ConvNets over more traditional segmentation algorithms on rodent MRI aligns with recent research (Roy et al., 2018; Liu et al., 2020; De Feo et al., 2021).

MedicDeepLabv3+ yielded the highest Dice coefficients, precision and recall, and the lowest HD (Table 1). Particularly, the outperformance of MedicDeepLabv3+ over the baseline DeepLabv3+ (Chen et al., 2018b) indicates that the proposed modifications (i.e., the incorporation of spatial attention layers and additional skip-connections), altogether,

led to improvements. Similar improvements after incorporating attention layers, such as the proposed spatial attention layers, have also been reported in the literature (Oktay et al., 2018; Wang et al., 2019; Tao et al., 2019; Xu et al., 2020). The same applies for adding skip connections (Drozdal et al., 2016; Li et al., 2018). In the brain midline area experiment, UNet achieved slightly higher Dice coefficients than the other 3D ConvNets. However, these Dice coefficients were computed only in the annotated slices, as finding the brain midline requires the manual annotations. As we showed in Table 1, Fig. 5, and the 17 Figures in Online Resource 1, 2D ConvNets, including UNet, produced misclassifications in the cerebellum and the olfactory bulb that were not annotated, leading to notably higher HD. Therefore, the small difference between UNet and the 3D ConvNets (Fig. 4) comes at the expense of those misclassifications that were disregarded during the evaluation. The differences among

Table 4 Dice, Hausdorff distance (HD), precision, and recall on the brain masks derived with MedicDeepLabv3+ in each cohort and time-point (TP) separately. Volumes indicate the number of volumes in the test set. Mean \pm std

Cohort	TP	Volumes	Dice	HD	Prec	Recall
1	2h	8	0.916 \pm 0.03	2.018 \pm 0.47	0.929 \pm 0.08	0.913 \pm 0.08
1	24h	8	0.912 \pm 0.10	2.581 \pm 1.82	0.894 \pm 0.18	0.952 \pm 0.02
2	24h	13	0.915 \pm 0.02	3.745 \pm 1.12	0.929 \pm 0.07	0.909 \pm 0.05
3	D35	16	0.949 \pm 0.02	2.142 \pm 0.44	0.956 \pm 0.03	0.942 \pm 0.02
4	24h	40	0.957 \pm 0.01	1.981 \pm 0.87	0.985 \pm 0.01	0.930 \pm 0.02
5	24h	23	0.947 \pm 0.01	1.557 \pm 0.43	0.935 \pm 0.03	0.960 \pm 0.02
6	D3	60	0.973 \pm 0.01	1.200 \pm 0.32	0.968 \pm 0.02	0.978 \pm 0.01
6	D28	58	0.946 \pm 0.02	1.688 \pm 0.60	0.937 \pm 0.04	0.957 \pm 0.02
7	D3	35	0.956 \pm 0.01	1.723 \pm 0.92	0.968 \pm 0.02	0.945 \pm 0.03
7	D21	35	0.950 \pm 0.01	1.548 \pm 0.75	0.937 \pm 0.02	0.964 \pm 0.01
8	24h	29	0.956 \pm 0.01	1.742 \pm 0.92	0.956 \pm 0.03	0.956 \pm 0.02
8	D3	26	0.953 \pm 0.01	1.599 \pm 0.35	0.949 \pm 0.02	0.957 \pm 0.01
8	D14	26	0.948 \pm 0.01	1.678 \pm 0.43	0.929 \pm 0.03	0.967 \pm 0.01
8	D28	23	0.943 \pm 0.02	1.713 \pm 0.36	0.940 \pm 0.03	0.946 \pm 0.03
9	24h	77	0.951 \pm 0.03	1.838 \pm 1.02	0.970 \pm 0.04	0.935 \pm 0.04
10	D7	36	0.919 \pm 0.05	2.226 \pm 0.81	0.880 \pm 0.10	0.970 \pm 0.02
11	24h	28	0.937 \pm 0.02	2.696 \pm 0.79	0.954 \pm 0.04	0.923 \pm 0.04
Average		541	0.948 \pm 0.03	1.833 \pm 0.89	0.949 \pm 0.05	0.951 \pm 0.03

Table 5 Dice, Hausdorff distance (HD), precision, and recall on the contralateral hemisphere masks derived with MedicDeepLabv3+ in each cohort and time-point (TP) separately. Volumes indicate the number of volumes in the test set. Mean \pm std

Cohort	TP	Volumes	Dice	HD	Prec	Recall
1	2h	8	0.883 \pm 0.03	3.593 \pm 0.83	0.871 \pm 0.06	0.904 \pm 0.07
1	24h	8	0.886 \pm 0.11	3.181 \pm 1.76	0.874 \pm 0.18	0.915 \pm 0.03
2	24h	13	0.876 \pm 0.03	3.792 \pm 1.60	0.902 \pm 0.08	0.859 \pm 0.05
3	D35	16	0.927 \pm 0.02	1.977 \pm 0.94	0.948 \pm 0.03	0.907 \pm 0.03
4	24h	40	0.928 \pm 0.02	1.889 \pm 0.76	0.962 \pm 0.03	0.898 \pm 0.03
5	24h	23	0.899 \pm 0.04	1.630 \pm 0.48	0.912 \pm 0.04	0.888 \pm 0.06
6	D3	60	0.951 \pm 0.02	2.766 \pm 2.27	0.936 \pm 0.04	0.967 \pm 0.01
6	D28	58	0.935 \pm 0.02	1.387 \pm 0.51	0.932 \pm 0.04	0.939 \pm 0.02
7	D3	35	0.930 \pm 0.02	2.038 \pm 1.40	0.959 \pm 0.02	0.904 \pm 0.04
7	D21	35	0.939 \pm 0.01	1.487 \pm 0.90	0.927 \pm 0.03	0.952 \pm 0.02
8	24h	29	0.935 \pm 0.02	2.574 \pm 2.32	0.936 \pm 0.03	0.936 \pm 0.03
8	D3	26	0.903 \pm 0.07	5.874 \pm 0.82	0.902 \pm 0.03	0.912 \pm 0.11
8	D14	26	0.933 \pm 0.02	1.911 \pm 1.34	0.901 \pm 0.04	0.967 \pm 0.01
8	D28	23	0.932 \pm 0.02	1.830 \pm 1.23	0.929 \pm 0.03	0.935 \pm 0.03
9	24h	77	0.917 \pm 0.03	2.234 \pm 1.67	0.926 \pm 0.05	0.911 \pm 0.04
10	D7	36	0.908 \pm 0.05	3.195 \pm 2.20	0.871 \pm 0.10	0.958 \pm 0.03
11	24h	28	0.894 \pm 0.03	4.200 \pm 2.27	0.900 \pm 0.04	0.892 \pm 0.04
Average		541	0.923 \pm 0.04	2.480 \pm 1.89	0.924 \pm 0.05	0.926 \pm 0.05

HighRes3DNet, V-Net, VoxResNet and MedicDeepLabv3+ were also very small. In contrast, the difference between MedicDeepLabv3+ and the baseline DeepLabv3+ was three times larger than between MedicDeepLabv3+ and UNet.

Our benchmark (Table 1) provides a valuable insight into whether 2D ConvNets produce better segmentations than 3D ConvNets on highly anisotropic data. In recent literature, 2D ConvNets appeared to be better (Jang et al., 2017; Isensee et al., 2017; Baumgartner et al., 2017), including in rodent images similar to our dataset (De Feo et al., 2021). 2D ConvNets outperformance may arise because contiguous slices can differ significantly in anisotropic data, thus, three-dimensional information might be unnecessary, and slice appearance might suffice to segment the regions of interest. Our data and, particularly, our manual annotations, were specially challenging since our regions of interest had similar intensity values to the cerebellum and olfactory bulb that were not annotated. Therefore, three-dimensional information can be critical to learn the location in the rostro-caudal axis of certain areas to avoid them. Indeed, our results support this intuition. Although Dice coefficient, precision and recall varied across architectures (Table 1), HD was consistently lower with 3D ConvNets. In other words, 2D ConvNets produced more critical misclassifications. Thus, our data showcased a scenario in which, despite the anisotropy, 3D ConvNets were superior to 2D ConvNets, showing that the architectural choices need to consider more specific information and not just whether the data is anisotropic.

We measured the discrepancy magnitude between the hemispheric ratio distributions from the segmentations and from the ground truth (Table 2), and V-Net and our

MedicDeepLabv3+ yielded the smallest effect size, indicating that the hemispheric ratios of their corresponding segmentations were more similar to the ground truth than the other ConvNets. We want to emphasize the importance of accurate hemispheric ratios as they are biomarkers for predicting acute stroke (Swanson et al., 1990; Gerriets et al., 2004). Both V-Net and MedicDeepLabv3+'s confidence intervals were zero-centered, and between these two, MedicDeepLabv3+'s was one third smaller than V-Net's. The effect size of VoxResNet and HighRes3DNet (the second and third best performing ConvNets after MedicDeepLabv3+) was much higher, and their confidence intervals did not include zero, which indicates that their hemispheric ratios were biased, being considerably larger than the ground truth. UNet's and DeepLabv3+'s effect size were also high, and DeepLabv3+'s confidence interval was the largest across all compared ConvNets. Thus, overall, and in agreement with the other experiments, MedicDeepLabv3+ compared favorably with the baseline DeepLabv3+ and the other competing methods.

ConvNets, and especially our MedicDeepLabv3+, produced segmentations more similar to the ground truth than the other methods (Table 1). Since these ConvNets were high capacity—requiring large GPU memory—and they were optimized with several images, their outperformance is in line with recent research (Tan & Le, 2019). However, annotated data are often scarce, and large GPU memory to optimize ConvNets is not necessarily available. Motivated by these constraints, we showed in two separate experiments that MedicDeepLabv3+ performed remarkably well with few annotated data and very limited GPU memory (see Tables 3,

4, and 5). In other words, our method can handle different scenarios without excessively sacrificing performance, which showcases MedicDeepLabv3+ generalization capabilities.

MedicDeepLabv3+ is publicly available, and it can be easily incorporated into existing pipelines, reducing human workload and accelerating rodent neuroimaging analyses. Furthermore, MedicDeepLabv3+ is fast, requires no pre-processing and postprocessing, and it can be optimized on MR images with different contrast, voxel resolution, field of view, lesion appearance, and limited GPU memory and annotated data. As hemisphere segmentation masks can be utilized in diverse studies, our work is relevant for multiple applications involving brain lesions in rat images.

Information Sharing Statement

The source code of MedicDeepLabv3+ and the trained models are publicly available under MIT License at <https://github.com/jmlipman/MedicDeepLabv3Plus>. All quantitative evaluation measures, to which conclusion in the work are based on, are presented in the Supplementary Materials. Training data may be available from Charles River Discovery Services upon the request of qualified parties but restrictions relating to the client privacy and intellectual property apply to the availability to these data, which were used under license for the current study. Typically, training data access will occur through collaboration and require interested parties to sign a material transfer agreement with Charles River Discovery Services prior to access to the training data.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s12021-022-09607-1>.

Acknowledgements The work of J.M. Valverde was funded from the European Union's Horizon 2020 Framework Programme (Marie Skłodowska Curie grant agreement #740264 (GENOMMED)). This work has also been supported by the grant #316258 from Academy of Finland (J. Tohka) and grant S21770 from the European Social Fund (R. De Feo). Part of the computational analysis was run on the servers provided by Bioinformatics Center, University of Eastern Finland, Finland.

Author Contributions Conceptualization: Juan Miguel Valverde, Artem Shatillo, Riccardo De Feo, Jussi Tohka; Data curation: Artem Shatillo; Formal Analysis: Juan Miguel Valverde, Jussi Tohka; Investigation: Juan Miguel Valverde; Funding acquisition: Jussi Tohka; Methodology: Juan Miguel Valverde, Riccardo De Feo, Jussi Tohka; Project administration: Jussi Tohka; Resources: Artem Shatillo; Software: Juan Miguel Valverde; Supervision: Jussi Tohka; Validation: Juan Miguel Valverde, Jussi Tohka; Visualization: Juan Miguel Valverde; Writing—original draft: Juan Miguel Valverde; Writing—review & editing: Juan Miguel Valverde, Artem Shatillo, Riccardo De Feo, Jussi Tohka.

Funding Open access funding provided by University of Eastern Finland (UEF) including Kuopio University Hospital. The work of J.M. Valverde was funded from the European Union's Horizon 2020 Framework Programme (Marie Skłodowska Curie grant agreement #740264 (GENOMMED)). This work has also been supported by the grant

#316258 from Academy of Finland (J. Tohka) and grant S21770 from the European Social Fund (R. De Feo).

Data Availability The trained models are available at <https://github.com/jmlipman/MedicDeepLabv3Plus>. Training data may be available from Charles River Discovery Services upon the request of qualified parties but restrictions relating to the client privacy and intellectual property apply to the availability to these data, which were used under license for the current study. Typically, training data access will occur through collaboration and require interested parties to sign a material transfer agreement with Charles River Discovery Services prior to access to the training data.

Code Availability The source code of MedicDeepLabv3+ and the trained models are publicly available at <https://github.com/jmlipman/MedicDeepLabv3Plus>.

Declarations

Ethics Approval All animal experiments were conducted according to the National Institute of Health (NIH) guidelines for the care and use of laboratory animals, and approved by the National Animal Experiment Board, Finland.

Consent to Participate Not applicable.

Consent for Publication Not applicable.

Conflicts of Interest/Competing Interests As disclosed in the affiliation section, Artem Shatillo, MD is a full-time payroll employee of the Charles River Discovery Services, Finland—a commercial preclinical contract research organization (CRO), which participated in the project and provided raw data as a part of companies R &D initiative.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Arnaud, A., Forbes, F., Coquery, N., Collomb, N., Lemasson, B., & Barbier, E. L. (2018). Fully automatic lesion localization and characterization: Application to brain tumors using multiparametric quantitative mri data. *IEEE Transactions on Medical Imaging*, 37, 1678–1689.
- Bae, M. H., Pan, R., Wu, T., & Badae, A. (2009). Automated segmentation of mouse brain images using extended mrf. *Neuroimage*, 46, 717–725.
- Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., Shinohara, R. T., Berger, C., Ha, S. M., Rozycki, M. et al. (2018). Identifying the best machine learning algorithms for brain tumor

- segmentation, progression assessment, and overall survival prediction in the brats challenge. arXiv preprint [arXiv:1811.02629](https://arxiv.org/abs/1811.02629)
- Baumgartner, C. F., Koch, L. M., Pollefeys, M., & Konukoglu, E. (2017). An exploration of 2d and 3d deep learning techniques for cardiac mr image segmentation. In *International Workshop on Statistical Atlases and Computational Models of the Heart* (pp. 111–119). Springer.
- Bernard, O., Lalonde, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.-A., Cetin, I., Lekadir, K., Camara, O., Ballester, M. A. G., et al. (2018). Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE Transactions on Medical Imaging*, *37*, 2514–2525.
- Carbone, L. (2021). Estimating mouse and rat use in american laboratories by extrapolation from animal welfare act-regulated species. *Scientific Reports*, *11*, 1–6.
- Chen, H., Dou, Q., Yu, L., Qin, J., & Heng, P.-A. (2018). Voxresnet: Deep voxelwise residual networks for brain segmentation from 3d mr images. *NeuroImage*, *170*, 446–455.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2014). Semantic image segmentation with deep convolutional nets and fully connected crfs. arXiv preprint [arXiv:1412.7062](https://arxiv.org/abs/1412.7062)
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2017). Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *40*, 834–848.
- Choi, C.-H., Yi, K. S., Lee, S.-R., Lee, Y., Jeon, C.-Y., Hwang, J., Lee, C., Choi, S. S., Lee, H. J., & Cha, S.-H. (2018a). A novel voxel-wise lesion segmentation technique on 3.0-t diffusion mri of hyperacute focal cerebral ischemia at 1 h after permanent mcao in rats. *Journal of Cerebral Blood Flow & Metabolism*, *38*, 1371–1383.
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., & Adam, H. (2018b). Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on Computer Vision (ECCV)* (pp. 801–818).
- Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1251–1258).
- Chou, N., Wu, J., Bingren, J. B., Qiu, A., & Chuang, K.-H. (2011). Robust automatic rodent brain extraction using 3-d pulse-coupled neural networks (pcnn). *IEEE Transactions on Image Processing*, *20*, 2554–2564.
- De Feo, R., Hmlinen, E., Manninen, E., Immonen, R., Valverde, J. M., Ndode-Ekane, X. E., Grhn, O., Pitknen, A., & Tohka, J. (2022). Convolutional neural networks enable robust automatic segmentation of the rat hippocampus in mri after traumatic brain injury. *Frontiers in Neurology*, *13*.
- De Feo, R., Shatillo, A., Sierra, A., Valverde, J. M., Gröhn, O., Giove, F., & Tohka, J. (2021). Automated joint skull-stripping and segmentation with multi-task u-net in large mouse brain mri databases. *NeuroImage*, (p. 117734).
- Dervieux, A., & Thomasset, F. (1980). A finite element method for the simulation of a rayleigh-taylor instability. In *Approximation methods for Navier-Stokes problems* (pp. 145–158). Springer.
- Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology*, *26*, 297–302.
- Dietterich, T. G. (2000). Ensemble methods in machine learning. In *International workshop on multiple classifier systems* (pp. 1–15). Springer.
- Drozdal, M., Vorontsov, E., Chartrand, G., Kadoury, S., & Pal, C. (2016). The importance of skip connections in biomedical image segmentation. In *Deep Learning and Data Labeling for Medical Applications* (pp. 179–187). Springer.
- Efron, B. (1987). Better bootstrap confidence intervals. *Journal of the American Statistical Association*, *82*, 171–185.
- Freret, T., Chazalviel, L., Roussel, S., Bernaudin, M., Schumann-Bard, P., & Boulouard, M. (2006). Long-term functional outcome following transient middle cerebral artery occlusion in the rat: correlation between brain damage and behavioral impairment. *Behavioral Neuroscience*, *120*, 1285.
- Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., & Lu, H. (2019). Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 3146–3154).
- Gerriets, T., Stolz, E., Walberer, M., Muller, C., Kluge, A., Bachmann, A., Fisher, M., Kaps, M., & Bachmann, G. (2004). Noninvasive quantification of brain edema and the space-occupying effect in rat stroke models using magnetic resonance imaging. *Stroke*, *35*, 566–571.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778).
- Heller, N., Isensee, F., Maier-Hein, K. H., Hou, X., Xie, C., Li, F., Nan, Y., Mu, G., Lin, Z., Han, M., et al. (2021). The state of the art in kidney and kidney tumor segmentation in contrast-enhanced ct imaging: Results of the kits19 challenge. *Medical Image Analysis*, *67*, 101821.
- Hsu, L.-M., Wang, S., Ranadive, P., Ban, W., Chao, T.-H. H., Song, S., Cerri, D. H., Walton, L. R., Broadwater, M. A., Lee, S.-H. et al. (2020). Automatic skull stripping of rat and mouse brain mri data using u-net. *Frontiers in Neuroscience*.
- Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint [arXiv:1502.03167](https://arxiv.org/abs/1502.03167)
- Isensee, F., Jaeger, P. F., Full, P. M., Wolf, I., Engelhardt, S., & Maier-Hein, K. H. (2017). Automatic cardiac disease assessment on cine-mri via time-series segmentation and domain specific features. In *International Workshop on Statistical Atlases and Computational Models of the Heart* (pp. 120–129). Springer.
- Isensee, F., Jaeger, P. F., Kohl, S. A., Petersen, J., & Maier-Hein, K. H. (2021). nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, *18*, 203–211.
- Jang, Y., Hong, Y., Ha, S., Kim, S., & Chang, H.-J. (2017). Automatic segmentation of lv and rv in cardiac mri. In *International Workshop on Statistical Atlases and Computational Models of the Heart* (pp. 161–169). Springer.
- Kervadek, H., Dolz, J., Tang, M., Granger, E., Boykov, Y., & Ayed, I. B. (2019). Constrained-cnn losses for weakly supervised segmentation. *Medical Image Analysis*, *54*, 88–99.
- Khan, Z., Yahya, N., Alsaih, K., Ali, S. S. A., & Meriaudeau, F. (2020). Evaluation of deep neural networks for semantic segmentation of prostate in t2w mri. *Sensors*, *20*, 3183.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *CoRR*, [abs/1412.6980](https://arxiv.org/abs/1412.6980).
- Koizumi, J., Yoshida, Y., Nakazawa, T., & Ooneda, G. (1986). Experimental studies of ischemic brain edema. 1. a new experimental model of cerebral embolism in rats in which recirculation can be introduced in the ischemic area. *Japanese Journal of Stroke*, *8*, 1–8.
- Kushibar, K., Valverde, S., Gonzalez-Villa, S., Bernal, J., Cabezas, M., Oliver, A., & Lladó, X. (2018). Automated sub-cortical brain structure segmentation combining spatial and deep convolutional features. *Medical Image Analysis*, *48*, 177–186.
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and anovas. *Frontiers in Psychology*, *4*, 863.
- Lee, C.-Y., Xie, S., Gallagher, P., Zhang, Z., & Tu, Z. (2015). Deeply-supervised nets. In *Artificial Intelligence and Statistics* (pp. 562–570). PMLR.

- Li, H., Xu, Z., Taylor, G., Studer, C., & Goldstein, T. (2018). Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems* (pp. 6389–6399).
- Li, W., Wang, G., Fidon, L., Ourselin, S., Cardoso, M. J., & Vercauteren, T. (2017). On the compactness, efficiency, and representation of 3d convolutional networks: brain parcellation as a pretext task. In *International Conference on Information Processing in Medical Imaging* (pp. 348–360). Springer.
- Liu, Y., Unsal, H. S., Tao, Y., & Zhang, N. (2020). Automatic brain extraction for rodent mri images. *Neuroinformatics*, (pp. 1–12).
- Ma, C., Ji, Z., & Gao, M. (2019). Neural style transfer improves 3d cardiovascular mr image segmentation on inconsistent data. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 128–136). Springer.
- Ma, D., Cardoso, M. J., Modat, M., Powell, N., Wells, J., Holmes, H., Wiseman, F., Tybulewicz, V., Fisher, E., Lythgoe, M. F., et al. (2014). Automatic structural parcellation of mouse brain mri using multi-atlas label fusion. *PLoS One*, *9*, e86576.
- McBride, D. W., Klebe, D., Tang, J., & Zhang, J. H. (2015). Correcting for brain swellings effects on infarct volume calculation after middle cerebral artery occlusion in rats. *Translational Stroke Research*, *6*, 323–338.
- Milletari, F., Navab, N., & Ahmadi, S.-A. (2016). V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)* (pp. 565–571). IEEE.
- Mulder, I. A., Khmelinskii, A., Dzyubachyk, O., de Jong, S., Rieff, N., Wermer, M. J., Hoehn, M., Lelieveldt, B. P., & van den Maagdenberg, A. M. (2017). Automated ischemic lesion segmentation in mri mouse brain data after transient middle cerebral artery occlusion. *Frontiers in Neuroinformatics*, *11*, 3.
- Murugavel, M., & Sullivan, J. M., Jr. (2009). Automatic cropping of mri rat brain volumes using pulse coupled neural networks. *NeuroImage*, *45*, 845–854.
- Myronenko, A., & Hatamizadeh, A. (2019). Robust semantic segmentation of brain tumor regions from 3d mris. In *International MIC-CAI Brainlesion Workshop* (pp. 82–89). Springer.
- Oguz, I., Zhang, H., Rumble, A., & Sonka, M. (2014). Rats: rapid automatic tissue segmentation in rodent brain mri. *Journal of Neuroscience Methods*, *221*, 175–182.
- Oktay, O., Schlemper, J., Folgoc, L. L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N. Y., Kainz, B. et al. (2018). Attention u-net: Learning where to look for the pancreas. arXiv preprint [arXiv:1804.03999](https://arxiv.org/abs/1804.03999)
- Osher, S., & Sethian, J. A. (1988). Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, *79*, 12–49.
- Pagani, M., Damiano, M., Galbusera, A., Tsiftaris, S. A., & Gozzi, A. (2016). Semi-automated registration-based anatomical labelling, voxel based morphometry and cortical thickness mapping of the mouse brain. *Journal of Neuroscience Methods*, *267*, 62–73.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L. et al. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* (pp. 8024–8035).
- 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* (pp. 234–241). Springer.
- Rote, G. (1991). Computing the minimum hausdorff distance between two point sets on a line under translation. *Information Processing Letters*, *38*, 123–127.
- Roy, S., Knutsen, A., Korotcov, A., Bosomtwi, A., Dardzinski, B., Butman, J. A., & Pham, D. L. (2018). A deep learning framework for brain extraction in humans and animals with traumatic brain injury. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)* (pp. 687–691). IEEE.
- Schwarz, A. J., Danckaert, A., Reese, T., Gozzi, A., Paxinos, G., Watson, C., Merlo-Pich, E. V., & Bifone, A. (2006). A stereotaxic mri template set for the rat brain with tissue class distribution maps and co-registered anatomical atlas: application to pharmacological mri. *NeuroImage*, *32*, 538–550.
- Swanson, R. A., Morton, M. T., Tsao-Wu, G., Savalos, R. A., Davidson, C., & Sharp, F. R. (1990). A semiautomated method for measuring brain infarct volume. *Journal of Cerebral Blood Flow & Metabolism*, *10*, 290–293.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1–9).
- Tan, M., & Le, Q. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning* (pp. 6105–6114). PMLR.
- Tao, Q., Ge, Z., Cai, J., Yin, J., & See, S. (2019). Improving deep lesion detection using 3d contextual and spatial attention. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 185–193). Springer.
- Valverde, J. M., Shatillo, A., De Feo, R., Gröhn, O., Sierra, A., & Tohka, J. (2020). Ratlesnetv2: A fully convolutional network for rodent brain lesion segmentation. *Frontiers in Neuroscience*, *14*, 1333.
- Wang, G., Shapey, J., Li, W., Dorent, R., Demetriadi, A., Bisdas, S., Paddick, I., Bradford, R., Zhang, S., Ourselin, S. et al. (2019). Automatic segmentation of vestibular schwannoma from t2-weighted mri by deep spatial attention with hardness-weighted loss. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 264–272). Springer.
- Wang, Y., Cheung, P.-T., Shen, G. X., Bhatia, I., Wu, E. X., Qiu, D., & Khong, P.-L. (2007). Comparing diffusion-weighted and t2-weighted mr imaging for the quantification of infarct size in a neonatal rat hypoxic-ischemic model at 24 h post-injury. *International Journal of Developmental Neuroscience*, *25*, 1–5.
- Wood, T. C., Lythgoe, D. J., & Williams, S. C. (2013). rbet: making bet work for rodent brains. In *Proceeding of the International Society for Magnetic Resonance in Medicine* (p. 2706). volume 21.
- Wu, T., Bae, M. H., Zhang, M., Pan, R., & Badea, A. (2012). A prior feature svm-mrf based method for mouse brain segmentation. *NeuroImage*, *59*, 2298–2306.
- Xie, Y., Lu, H., Zhang, J., Shen, C., & Xia, Y. (2019). Deep segmentation-emendation model for gland instance segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 469–477). Springer.
- Xu, X., Lian, C., Wang, S., Wang, A., Royce, T., Chen, R., Lian, J., & Shen, D. (2020). Asymmetrical multi-task attention u-net for the segmentation of prostate bed in ct image. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 470–479). Springer.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.