



# RatLesNetv2: A Fully Convolutional Network for Rodent Brain Lesion Segmentation

Juan Miguel Valverde<sup>1\*</sup>, Artem Shatillo<sup>2</sup>, Riccardo De Feo<sup>1,3,4</sup>, Olli Gröhn<sup>1</sup>, Alejandra Sierra<sup>1</sup> and Jussi Tohka<sup>1</sup>

<sup>1</sup>A.I. Virtanen Institute for Molecular Sciences, University of Eastern Finland, Kuopio, Finland, <sup>2</sup>Charles River Discovery Services, Kuopio, Finland, <sup>3</sup>Centro Fermi-Museo Storico della Fisica e Centro Studi e Ricerche Enrico Fermi, Rome, Italy, <sup>4</sup>Sapienza Università di Roma, Rome, Italy

## OPEN ACCESS

### Edited by:

Tim B. Dyrby,  
Technical University of Denmark,  
Denmark

### Reviewed by:

Yi Zhang,  
Zhejiang University, China  
Shanshan Jiang,  
Johns Hopkins Medicine,  
United States

### \*Correspondence:

Juan Miguel Valverde  
juanmiguel.valverde@uef.fi

### Specialty section:

This article was submitted to  
Brain Imaging Methods,  
a section of the journal  
Frontiers in Neuroscience

**Received:** 25 September 2020

**Accepted:** 25 November 2020

**Published:** 22 December 2020

### Citation:

Valverde JM, Shatillo A, De Feo R, Gröhn O, Sierra A and Tohka J (2020) RatLesNetv2: A Fully Convolutional Network for Rodent Brain Lesion Segmentation. *Front. Neurosci.* 14:610239. doi: 10.3389/fnins.2020.610239

We present a fully convolutional neural network (ConvNet), named RatLesNetv2, for segmenting lesions in rodent magnetic resonance (MR) brain images. RatLesNetv2 architecture resembles an autoencoder and it incorporates residual blocks that facilitate its optimization. RatLesNetv2 is trained end to end on three-dimensional images and it requires no preprocessing. We evaluated RatLesNetv2 on an exceptionally large dataset composed of 916 T2-weighted rat brain MRI scans of 671 rats at nine different lesion stages that were used to study focal cerebral ischemia for drug development. In addition, we compared its performance with three other ConvNets specifically designed for medical image segmentation. RatLesNetv2 obtained similar to higher Dice coefficient values than the other ConvNets and it produced much more realistic and compact segmentations with notably fewer holes and lower Hausdorff distance. The Dice scores of RatLesNetv2 segmentations also exceeded inter-rater agreement of manual segmentations. In conclusion, RatLesNetv2 could be used for automated lesion segmentation, reducing human workload and improving reproducibility. RatLesNetv2 is publicly available at <https://github.com/jmlipman/RatLesNetv2>.

**Keywords:** ischemic stroke, lesion segmentation, deep learning, rat brain, magnetic resonance imaging

## 1. INTRODUCTION

Rodents frequently serve as models for human brain diseases. They account for more than 80% of the animals used in research in recent years (Dutta and Sengupta, 2016). In addition to basic research, rodent models are important in, for example, drug discovery and the development of new treatments. *In vivo* imaging of rodents is used for monitoring disease progression and therapeutic response in longitudinal studies. In particular, magnetic resonance imaging (MRI) is essential in pre-clinical studies for conducting quantitative analyses due to its non-invasiveness and versatility. As an example, the quantification of brain lesions requires segmenting the lesions, and the lack of reliable tools to automate rodent brain lesion segmentation forces researchers to segment these images manually.

Manual segmentation can be prohibitively time-consuming as studies involving animals may acquire hundreds of three-dimensional (3D) images. Furthermore, the difficulty of defining lesion boundaries leads to moderate inter- and intra-rater agreement; previous studies have reported that Dice coefficients (Dice, 1945) between annotations made by two humans can be as low as 0.73

(Valverde et al., 2019) or 0.79 (Mulder et al., 2017a). Moderate inter-rater agreement is caused by several factors that affect the segmentation quality, including partial volume effect, image contrast and annotator's knowledge and experience. Despite these liabilities, manual segmentation is the gold standard and a common practice among researchers who use animal models (Moraga et al., 2016; De Feo and Giove, 2019).

Semi-automatic methods are a faster alternative to manual segmentation. However, they fail to overcome the subjectivity of the manual segmentation, as human interaction is required. To the best of the authors' knowledge, there are only two studies that introduce and evaluate semi-automatic algorithms for rodent brain lesion segmentation. Wang et al. (2007) evaluated a combination of thresholding operations commonly used in the literature to segment lesions on apparent diffusion coefficient (ADC) maps and T2-weighted images. Choi et al. (2018) first normalized the intensity values of each image with respect to the contralateral hemisphere of the brain, and they performed a series of thresholding operations to segment permanent middle cerebral artery occlusion ischemic lesions in 31 diffusion-weighted images (DWIs) of the rat brain. Both methods require the manual segmentation of the contralateral hemisphere. Additionally, these thresholding-based and other voxel-wise approaches disregard the spatial and contextual information of the images, and they are sensitive to the image modality, contrast, and possible artifacts. Pipelines that rely on thresholding operations may result in poor and inconsistent segmentation results in the form of holes within and outside the lesion mask (Figure 1).

For lesion segmentation in rodent MRI, researchers have proposed a few fully-automated methods in recent years. Mulder et al. (2017a) developed a level-set-based algorithm that was tested on 121 T2-weighted mouse brain scans. However, the accuracy of their method heavily relies on the performance of other independent steps, such as registration, skull-stripping and contralateral ventricle segmentation. Arnaud et al. (2018) derived a pipeline that detects voxels that are anomalous with respect to a reference model of healthy animals, and they evaluated the pipeline on 53 rat brain MRI maps. Nonetheless, this pipeline

was specifically designed for quantitative MRI, and it expects sham-operated animals in the data set, a requirement that is not always feasible.

Deep learning, and more specifically convolutional neural networks (ConvNets), has become increasingly popular due to its competitive performance in medical image segmentation. Literature on brain lesion segmentation in MR images with ConvNets is dominated by approaches tested on human-derived data (e.g., Duong et al., 2019; Gabr et al., 2019; Yang et al., 2019). Despite using ConvNets, typical brain lesion segmentation approaches are multi-step, i.e., they rely on preprocessing procedures, such as noise reduction, registration, skull-stripping and inhomogeneity correction. Therefore, the performance of the preprocessing steps influences the quality of the final segmentation. In contrast to human-derived data, rodent segmentation data sets are scarce and smaller in size (Mulder et al., 2017b); consequently, ConvNet-based segmentation methods benchmarked on rodent MR images are rare. An exception—not in the lesion segmentation—is Roy et al. (2018)'s work, which introduced a framework to extract brain tissue (i.e., skull-stripping) on human and mice MRI scans after traumatic brain injury.

We present RatLesNetv2, the first 3D ConvNet for segmenting rodent brain lesions in pre-clinical MR images. Our fully-automatic approach is trained end to end, requires no preprocessing, and it was validated on a large and diverse data set composed by 916 MRI rat brain scans at nine different lesion stages from 671 rats utilized to study focal cerebral ischemia. We extend our earlier conference paper (Valverde et al., 2019) by (1) improving our previous ConvNet (Valverde et al., 2019) with a deeper and different architecture and providing an ablation study (Meyes et al., 2019) justifying certain architectural choices; (2) evaluating the generalization capability of our model on a considerably larger and more heterogeneous data set via Dice coefficient, compactness and Hausdorff distance under different training settings (training set size and different ground truth); and (3) making RatLesNetv2 publicly available.

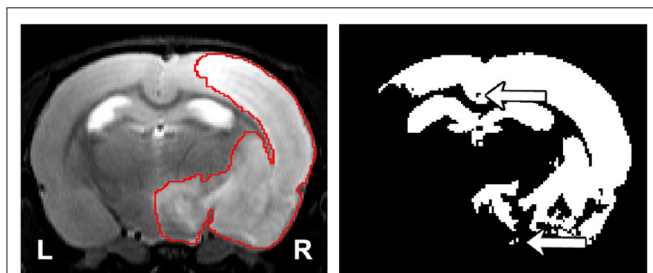
We show that RatLesNetv2 generates more realistic segmentations than our previous RatLesNet, and than 3D U-Net (Çiçek et al., 2016) and VoxResNet (Chen et al., 2018a), two state-of-the-art ConvNets specifically designed for medical image segmentation. Additionally, the Dice coefficients of the segmentations derived with RatLesNetv2 exceeded inter-rater agreement scores.

## 2. MATERIALS AND METHODS

### 2.1. Data

The data set consisted of 916 MR T2-weighted brain scans of 671 adult male Wistar rats weighting between 250 and 300 g. The data, provided by Discovery Services site of Charles River Laboratories,<sup>1</sup> were derived from 12 different studies. Transient (120 min) focal cerebral ischemia was produced by middle cerebral artery occlusion in the right hemisphere of the brain (Koizumi et al., 1986). MR data acquisitions were

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



**FIGURE 1** | (Left) Representative lesion with its ground truth. (Right) Segmentation of the lesion using thresholding where the threshold was found by maximizing the Dice coefficient with respect to the manual segmentation. The arrows indicate the presence of holes and islands (independently connected components) within and outside the mask, respectively. The hippocampus and ventricles were entirely misclassified as lesion.

**TABLE 1** | Number of scans per study segregated by lesion stage, including sham-operated animals.

Study	2 h	24 h	D3	D7	D14	D21	D28	D35	Shams
A	12	12	0	0	0	0	0	0	24
B	0	46	0	0	0	0	0	0	3
C	0	59	0	0	0	0	0	0	1
D	0	162	0	0	0	0	0	0	4
E	0	0	0	0	0	0	0	20	1
F	0	33	30	0	30	0	27	0	46
G	0	0	0	53	0	0	0	0	12
H	0	45	0	0	0	0	0	0	0
I	0	0	64	0	0	0	62	0	0
J	0	32	0	0	0	0	0	0	0
K	0	17	0	0	0	0	0	0	0
L	0	0	41	0	0	40	0	0	40
Total	12	406	135	53	30	40	89	20	131

performed at different time-points after the occlusion (for details, see **Table 1**). Some studies also had sham-operated animals that underwent identical surgical procedures, but without the actual occlusion. 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  $30 \times 30 \text{ mm}^2$  producing  $256 \times 256$  imaging matrices of resolution  $117 \times 117 \mu\text{m}$ . No MRI preprocessing steps, such as inhomogeneity correction, artifact removal, registration or skull stripping, were applied to the T2-weighted images. Images were zero-centered and their variance was normalized to one.

The provided lesion segmentations were annotated by several trained technicians employed by Charles River. We performed an additional independent manual segmentation of the lesions on the first study that was acquired (study A, **Table 1**) to approximate inter-rater variability. The average Dice coefficient (Dice, 1945) between the two manual segmentations was 0.67 with a standard deviation of 0.12 on 2 h lesions and 0.79 with a standard deviation of 0.08 on 24 h lesions. The overall average was  $0.73 \pm 0.12$ . Unless stated otherwise, we used our independent segmentation as the ground truth for study A.

We produced two different train/test set divisions. (1) In the first one, the training set contained the 48 scans of the study which was used to approximate inter-rater variability (study A, **Table 1**) and the test set contained the remaining 868 images. The training set was further divided to training (36 images) and validation sets (12 images). This train/test division is referred to as “homogeneous” and its train/validation split has the same ratio 2/24 h time-points and sham/no-sham animals. (2) The second division also contained 48 training scans and the test set contained 868 scans, but the training set was different

from the homogeneous division. This division is referred to as “heterogeneous” because the training set was more diverse. The training set was divided into training (40 images) and validation (8 images) set. The training and the validation sets were formed by 5 and 1 images per lesion time-point, respectively, with no images from sham-operated animals. The size of our training set was deliberately much smaller than the test set for two reasons: (1) to replicate the typical pre-clinical setting in which rodent MR images are few and (2) to create a large and representative test set.

## 2.2. Convolutional Neural Networks

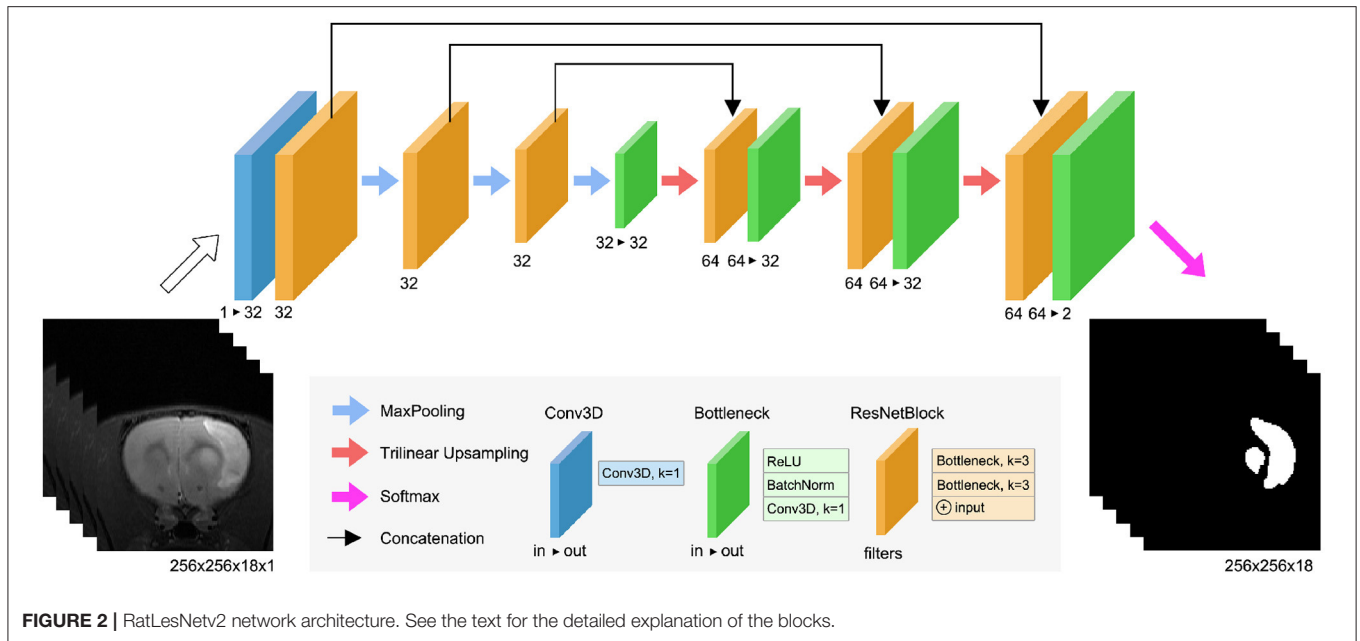
Convolutional neural networks (ConvNets) use stacks of convolutions to transform spatially correlated data, such as images, to extract their features. The first layers of the network capture low-level information, such as edges and corners, and the final layers extract more abstract features. The number of convolutions adjusts two attributes of ConvNets: parameter number and network depth. An excessive number of parameters leads to overfitting—memorizing the training data; an insufficient number of parameters constrains the learning capability of the model. Model depth is associated with the number of times the input data is transformed, and this depth also adjusts the area that influences the prediction—the receptive field (RF). Recent approaches reduce model parameters while maintaining the RF by using more stacked convolutions of smaller kernel size (Szegedy et al., 2016).

Model architectures based on U-Net (Ronneberger et al., 2015) are popular in medical image segmentation tasks. In contrast to patch-based models, the input images and the generated masks are the same size, which makes U-Nets computationally more efficient to train and to evaluate. The U-Net architecture resembles an autoencoder with skip connections between the same levels of the encoder and decoder. The encoder transforms and reduces the dimensionality of the input images, and the decoder recovers the spatial information with the help of skip connections.

Skip connections also facilitate the gradient flow during back-propagation (Drozdal et al., 2016), but they are not sufficient to prevent the gradient of the loss to vanish, which makes the network harder to train. This is also referred as the vanishing gradient problem (He et al., 2016), and it particularly affects the final layers of the encoder part. Adding residual connections (He et al., 2016) along the network alleviates the vanishing gradient problem and it also yields in faster convergence rates during the optimization (Drozdal et al., 2016).

## 2.3. RatLesNetv2 Architecture

RatLesNetv2 (**Figure 2**) has three downsampling and three upsampling stages connected via skip connections. Maxpooling downsamples the data with a window size and strides of 2, and trilinear interpolation upsamples the feature maps. Bottleneck layers (**Figure 2**, green blocks) stack a ReLU activation function, a batch normalization (BatchNorm) layer (Ioffe and Szegedy, 2015) and a 3D convolution with kernel size of 1 that combines and modifies the number of channels of the feature maps from *in* to *out*. ResNetBlock layers (**Figure 2**, orange blocks) contain two stacks of ReLU activations, BatchNorm, and 3D convolutions



**FIGURE 2** | RatLesNetv2 network architecture. See the text for the detailed explanation of the blocks.

with kernel size of 3. Similarly to VoxResNet (Chen et al., 2018a), the input and output of each block is summed in a ResNet-style (He et al., 2016). The width of the blocks in the decoder is twice (64) with respect to the encoder part (32) due to the concatenation of previous layers in the same stage of the network.

At the end of the network, the probabilities  $\mathbf{z} = [z_1, z_2]$  (corresponding to non-lesion and lesion labels) for each voxel are normalized by the Softmax function

$$q_i = \text{Softmax}(\mathbf{z})_i = \frac{e^{z_i}}{\sum_{j=1}^2 e^{z_j}}, \quad (1)$$

and the segmentation label is  $\arg \max_i(q_i), i = 1, 2$ .

RatLesNetv2 architecture differs from our previous RatLesNet (Valverde et al., 2019) in two aspects. First, RatLesNetv2 has one additional downsampling and upsampling level, increasing the receptive field to  $76 \times 76 \times 76$  voxels. These extra levels allows RatLesNetv2 to consider more information from a larger volume. Second, RatLesNetv2 replaces unpooling (Noh et al., 2015) and DenseNetBlocks (Huang et al., 2017) with trilinear upsampling and ResNetBlocks, respectively, reducing memory usage and execution time. In contrast to VoxResNet (Chen et al., 2018a), RatLesNetv2 architecture resembles an autoencoder, and RatLesNetv2 employs no transposed convolutions, reducing the number of parameters. Additionally, unlike 3D U-Net (Çiçek et al., 2016), RatLesNetv2 uses residual blocks that reuse previous computed feature maps and facilitate the optimization.

## 2.4. Loss Function

ConvNets’ parameters are optimized by minimizing a loss function that describes the difference between the predictions and the ground truth. RatLesNetv2 was optimized with Adam (Kingma and Ba, 2014) by minimizing cross entropy and Dice loss functions  $L_{total} = L_{BCE} + L_{Dice}$ . Cross entropy measures

the error as the difference between distributions. Since our annotations consist of only two classes (lesion and non-lesion) we used binary cross entropy

$$L_{BCE} = -\frac{1}{N} \sum_{i=1}^N p_i \cdot \log(q_i) + (1 - p_i) \cdot \log(1 - q_i), \quad (2)$$

where  $p_i \in \{0, 1\}$  represents whether voxel  $i$  is lesion in the ground truth and  $q_i \in [0, 1]$  is the predicted Softmax probability of lesion class. Dice loss (Milletari et al., 2016) is defined as:

$$L_{Dice} = 1 - \frac{2 \sum_i^N p_i q_i}{\sum_i^N p_i^2 + \sum_i^N q_i^2}. \quad (3)$$

The rationale behind using Dice loss is to directly maximize the Dice coefficient, one of the metrics to assess image segmentation performance. Although the derivative of Dice loss can be unstable when its denominator is very small, the use of BatchNorm and skip connections helps during the optimization by smoothing the loss landscape (Li et al., 2018; Santurkar et al., 2018).

## 2.5. Post-processing

Since our model optimizes a per-voxel loss function, small undesirable clusters of voxels may appear disconnected from the main predicted mask. These spurious clusters may be referred as “islands” when they are separated from the largest connected component and “holes” when they are inside the lesion mask. Figure 1 illustrates these terms.

Small islands and holes can be removed in a final post-processing operation, yielding more realistic segmentations. Determining the maximum size of these holes and islands is, however, challenging in practice: A very small threshold will not eliminate enough small islands and a too large threshold may

remove small lesions. In our experiments, we chose a threshold such that 90% of the holes and islands in the training data were removed. More specifically, we removed holes and islands of 20 voxels or less, inside and outside the lesion masks.

## 2.6. Evaluation Metrics

We assessed the performance of each ConvNet by measuring the Dice coefficient, Hausdorff distance and compactness. In agreement with the literature (Fenster and Chiu, 2005), we argue that Dice coefficient alone is not an effective measure in rodent lesion segmentation, which is why we complemented it with the two other metrics.

### 2.6.1. Dice Coefficient

Dice coefficient (Dice, 1945) is one of the most popular metrics in the field of image segmentation. It measures the overlap volume between two binary masks, typically the prediction of the model and the manually-annotated ground truth. Dice coefficient is formally described as:

$$Dice(A, B) = \frac{2|A \cap B|}{|A| + |B|}, \quad (4)$$

where  $A$  and  $B$  are the segmentation masks.

### 2.6.2. Compactness

Compact lesion masks are realistic and resemble human-made annotations. Compactness can be defined as the ratio between surface area (*area*) and volume of the mask (*volume*) (Bribiesca, 2008). More specifically, we define compactness as:

$$Compactness = area^{1.5} / volume, \quad (5)$$

which has a constant minimum value of  $6\sqrt{\pi}$  for any sphere. Compactness measure penalizes holes, islands and non-smooth borders because these increase the surface area with respect to the volume. Therefore, low compactness values that describe compact segmentations are desirable.

### 2.6.3. Hausdorff Distance

Hausdorff distance (HD) (Rote, 1991) is defined as:

$$d(A, B) = \max \left\{ \max_{a \in \partial A} \min_{b \in \partial B} |b - a|, \max_{b \in \partial B} \min_{a \in \partial A} |a - b| \right\}, \quad (6)$$

where  $A$  and  $B$  are the segmentation masks, and  $\partial A$  and  $\partial B$  are their respective boundary voxels. It measures the maximum distance of the ground truth surface to the closest voxel of the prediction, i.e., the largest segmentation error. Measuring Hausdorff distance in brain lesion segmentation studies is crucial since misclassifications far from the lesion boundaries are more severe. The reported Hausdorff distances were in millimeters.

Hausdorff distance and compactness values were calculated exclusively in animals with lesions. Hausdorff distance values on slightly imperfect segmentations of sham-operated animals are excessively large and distort the overall statistics. Additionally, compactness can not be calculated on empty volumes derived from scans without lesions. Voxel anisotropy was accounted

for when computing HD and compactness. Finally, we assessed significance of performance difference through a paired permutation test with 10,000 random iterations on the post-processed segmentations with 0.05 as the significance threshold.

## 2.7. Experimental Setup

### 2.7.1. Training

RatLesNetv2, 3D U-Net (Çiçek et al., 2016), VoxResNet (Chen et al., 2018a) and RatLesNet (Valverde et al., 2019) were optimized with Adam (Kingma and Ba, 2014) ( $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-8}$ ), starting with a learning rate of  $10^{-5}$  for 700 epochs. A small set of learning rates were tested on each architecture to ensure that we used the best performing learning rate in each model. Models were randomly initialized and trained three times separately, and their performance was evaluated from the lesion masks derived with majority voting across these three independent runs. In other words, for each architecture we ensembled three independently trained models. We confirmed that this strategy, typical to remove uncorrelated errors (Dietterich, 2000), improves performance.

### 2.7.2. Experiments

#### 2.7.2.1. Performance Comparison

We optimized RatLesNetv2, 3D U-Net (Çiçek et al., 2016), VoxResNet (Chen et al., 2018a) and RatLesNet (Valverde et al., 2019) on both the homogeneous and heterogeneous data set divisions (section 2.1) and compared their performance.

#### 2.7.2.2. Ablation Study

We conducted an ablation study (Meyes et al., 2019) in which we changed or removed certain parts of the model to comprehend the effects of the characteristics of RatLesNetv2 architecture. More specifically, we modified the interconnections between layers within each block, changed the number of downsampling/upsampling blocks, and increased and decreased the number of filters.

#### 2.7.2.3. Ground Truth Disparity Effect

We trained two separate RatLesNetv2 models on segmentations annotated by two different operators. This can be seen as an inter-rater variability study of the same ConvNet with disparate knowledge. We run RatLesNetv2 three times for each ground truth on the homogeneous training data, which come exclusively from the study with the two annotations (Study A, Table 1). RatLesNetv2 produced six sets of 868 masks  $\hat{y}_{g,r}$  where  $g \in \{1, 2\}$  refers to the annotator segmenting the training data and  $r \in \{1, 2, 3\}$  refers to the run. First, we approximated the intra-rater variability of RatLesNetv2 by calculating the Dice coefficients among the three runs for each ground truth separately, i.e.,  $\{dice(\hat{y}_{g,1}, \hat{y}_{g,2}), dice(\hat{y}_{g,2}, \hat{y}_{g,3}), dice(\hat{y}_{g,1}, \hat{y}_{g,3})\}$  for  $g = 1, 2$ . This led to two sets of three Dice coefficients per mask. Second, we calculated the Dice coefficient of the masks across the different ground truths  $\{dice(\hat{y}_{1,i}, \hat{y}_{2,j})\}$  for  $i, j = 1, 2, 3$  to approximate inter-rater similarity, leading to nine Dice coefficients per mask.

#### 2.7.2.4. Training Set Size

We optimized RatLesNetv2 with training sets of different sizes to understand the relation between training set size and

generalization capability. The training sets had the same ratio of time-points, i.e., we enlarged the training sets by 1 sample per time-point. Since the lowest number of samples across time-points corresponds to 12 (2 h lesions) and we want to keep at least 1 image per time-point in the test set, we produced 11 training sets  $T_i$  of size  $|T_i| = 8i$  for  $i = 1, \dots, 11$ , where 8 is the number of lesion stages.

### 2.7.3. Implementation

RatLesNetv2 was implemented in Pytorch (Paszke et al., 2019) and it was run on Ubuntu 16.04 with an Intel Xeon W-2125 CPU @ 4.00 GHz processor, 64 GB of memory and an NVidia GeForce GTX 1080 Ti with 11 GB of memory. RatLesNetv2 is publicly available at <https://github.com/jmlipman/RatLesNetv2>.

## 3. RESULTS

### 3.1. Performance of RatLesNetv2

Table 2 lists the quantitative validation results on the test set excluding sham-operated animals that typically yield Dice coefficients of 1.0. As can be seen in Table 2, RatLesNetv2 produced similar or better Dice coefficients and Hausdorff distances, and more compact segmentations than the other ConvNets. The average Dice coefficients varied from 0.784 (homogeneous division) to 0.813 (heterogeneous division). Dice coefficients had a large standard deviation regardless of the architecture (from 0.15 to 0.20). However, note that the sample-wise difference between the Dice coefficients of RatLesNetv2 and VoxResNet had a smaller standard deviation of 0.05, i.e., the Dice values between different networks were correlated. Table 2 shows that RatLesNetv2 achieved significantly better compactness values (all  $p$ -values  $< 0.011$ ) than 3D U-Net, VoxResNet and RatLesNet. Remarkably, 3D U-Net and VoxResNet produced masks with non-smooth borders and several more holes and islands, leading to less compact segmentations (see Figure 3 and Figures in the Supplementary Material). The average compactness values of RatLesNetv2 were higher than the ground truth ( $20.98 \pm 3.28$ ,  $p = 0.003$ ); this was expected as human annotators are likely to produce segmentations with excessively rounded boundaries.

Post-processing had little to no effect on the average Dice coefficients, but it enhanced the final segmentation quality as it removed spurious clusters of voxels. This improvement was reflected in the reduction of compactness values and the considerable decrease of Hausdorff distances. Remarkably, the difference in the Hausdorff distances before and after post-processing was more pronounced in 3D U-Net, VoxResNet and RatLesNet.

Table 3 lists the quantitative results by lesion stage to understand the performance of RatLesNetv2 in detail. Training RatLesNetv2 on the homogeneous data division, whose training set included almost twice as many 24 h lesion scans as the heterogeneous division (9 scans vs. 5 scans), led to a slight increase in the average Dice coefficient and Hausdorff distance in 24 h lesion scans. However, there was no significant difference between either the Dice coefficients ( $p = 0.057$ ) nor Hausdorff distances ( $p = 0.08$ ) of the segmentations derived in the

**TABLE 2 |** Performance evaluation on the test set before and after post-processing.

Model	Dice (no shams)	Compactness	HD
RatLesNetv2-post	<b>0.784 ± 0.18<sup>a</sup></b>	<b>29.332 ± 7.86<sup>b</sup></b>	3.522 ± 3.64
RatLesNetv2	0.784 ± 0.18	29.609 ± 8.12	3.687 ± 3.30
3D U-Net-post	0.769 ± 0.20	36.741 ± 11.41	3.665 ± 3.81
3D U-Net	0.768 ± 0.20	37.599 ± 11.77	4.097 ± 3.69
VoxResNet-post	0.757 ± 0.19	37.096 ± 13.00	3.692 ± 3.46
VoxResNet	0.757 ± 0.19	38.161 ± 13.62	4.943 ± 3.38
RatLesNet-post	0.742 ± 0.18	35.045 ± 10.71	3.892 ± 2.54
RatLesNet	0.741 ± 0.18	35.888 ± 10.76	4.679 ± 2.55
RatLesNetv2-post	0.813 ± 0.16	<b>23.105 ± 4.58<sup>c</sup></b>	3.334 ± 3.34
RatLesNetv2	0.813 ± 0.16	23.177 ± 4.64	3.512 ± 3.31
3D U-Net-post	0.813 ± 0.15	28.247 ± 5.92	3.099 ± 2.47
3D U-Net	0.812 ± 0.15	28.639 ± 5.99	3.221 ± 2.47
VoxResNet-post	0.806 ± 0.14	32.937 ± 10.05	3.585 ± 3.27
VoxResNet	0.805 ± 0.14	33.634 ± 10.53	4.535 ± 3.46
RatLesNet-post	0.764 ± 0.15	31.348 ± 9.66	3.218 ± 2.79
RatLesNet	0.764 ± 0.15	31.669 ± 9.86	3.384 ± 2.56

Average Dice coefficients were reported in images of animals with lesions. Top: Homogeneous division. Bottom: Heterogeneous division. Bold: Values significantly better than the other architectures (<sup>a</sup> $p = 0.007$ , <sup>b</sup> $p = 0.011$ , <sup>c</sup> $p = 0.005$ ).

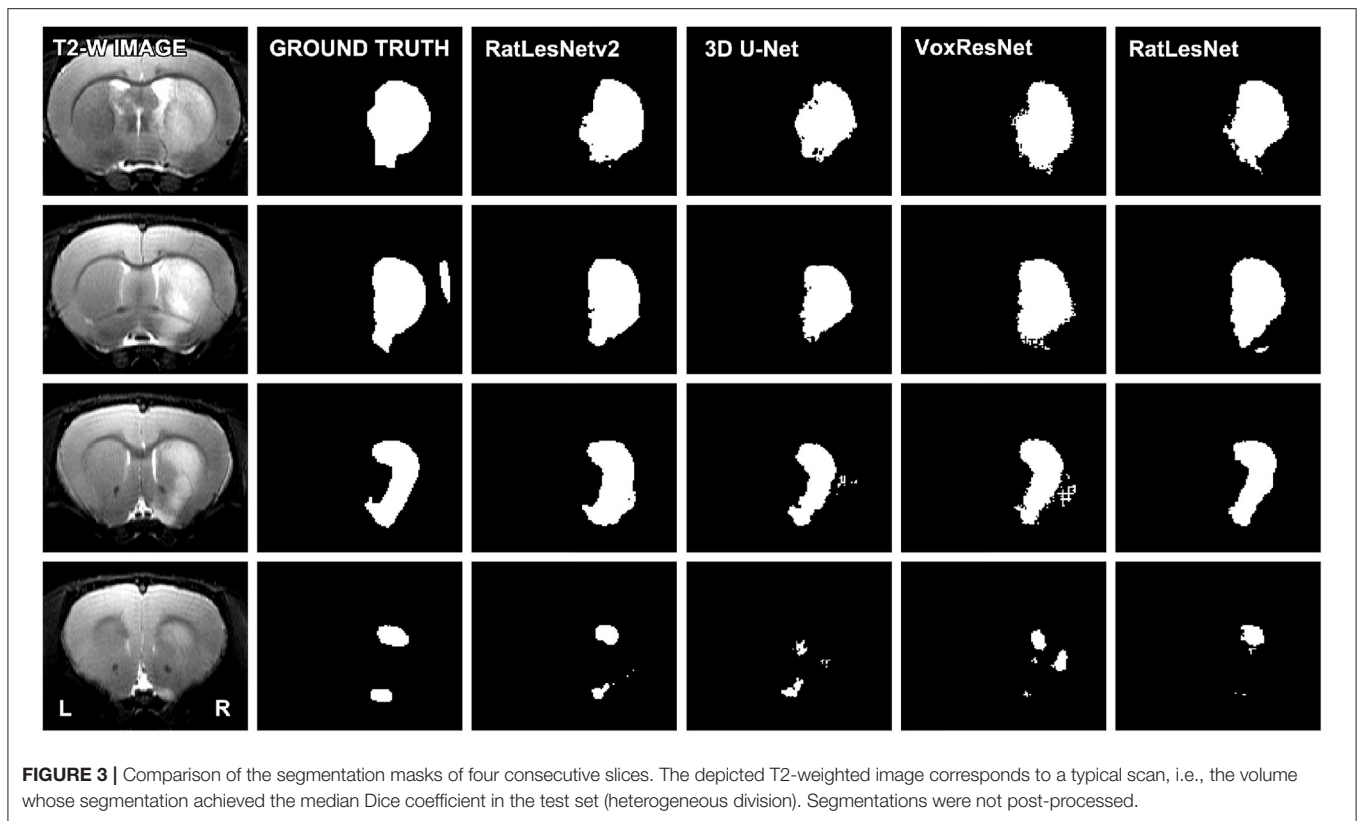
two cases. Dice coefficients, compactness values and Hausdorff distances of the segmentations produced after training on the homogeneous division deteriorated as the time-point was farther from 2 and 24 h.

Training on the heterogeneous training set notably improved the average Dice coefficients and compactness values of every model (Table 2) and every time-point (Table 3) with respect to homogeneous division, except on 24 h lesions. Furthermore, it decreased the standard deviation of the Dice coefficients and compactness values. RatLesNetv2 recognized animals without lesions notably well even if they were not part of the training set, providing average Dice coefficients of 1.0 on sham-operated animals even without post-processing. Additionally, Dice coefficients on 2 h lesions, 24 h lesions, and overall were higher than inter-rater agreement.

Ensembling three ConvNets of the same architecture optimized on the same training set led to significantly better performance scores in all cases (all  $p$ -values  $< 0.007$ ) as it discarded small segmentation inconsistencies. This strategy increased Dice coefficients by an average of 2% and decreased compactness and Hausdorff distances by an average of 5 and 23% with respect to the first run. The Dice coefficients, compactness values and Hausdorff distances from the individual images used for calculating the reported statistics are also included in the Supplementary Materials as CSV files.

### 3.2. Ablation Studies

The performance scores of RatLesNetv2 after modifying its architecture during the ablation studies are reported in Table 4.



**FIGURE 3** | Comparison of the segmentation masks of four consecutive slices. The depicted T2-weighted image corresponds to a typical scan, i.e., the volume whose segmentation achieved the median Dice coefficient in the test set (heterogeneous division). Segmentations were not post-processed.

### 3.2.1. DenseNetBlock

Similarly to RatLesNet (Valverde et al., 2019), DenseNet-style (Huang et al., 2017) blocks were implemented in RatLesNetv2 while keeping the same number of parameters of the baseline RatLesNetv2 model. Dice coefficients and compactness values were significantly deteriorated with respect to RatLesNetv2 baseline (all  $p$ -values  $< 0.037$ ), and Hausdorff distances increased slightly in homogeneous data division, whereas they decreased in heterogeneous division. Additionally, DenseNetBlocks demanded notably more memory due to the concatenation operation.

### 3.2.2. Halving the Receptive Field (RF)

The third downsampling stage of RatLesNetv2 was eliminated in order to reduce the receptive field from 72 voxels down to 36. An additional test (marked in Table 4 with an \*) matched the number of parameters to the baseline. The reduction of the receptive field yielded in significant improvements of the Dice coefficient and a significant deterioration of the compactness and Hausdorff distance in the heterogeneous division (all  $p$ -values  $< 0.028$ ). On the other hand, in the homogeneous division Dice coefficients and compactness values were worse than RatLesNetv2 baseline.

### 3.2.3. Network Width

We increased and decreased the number of filters of RatLesNetv2 by 4 (Table 4, Width-28 and Width-36). This modification decreased the Dice coefficients with respect to RatLesNetv2 and

led to no significant difference in the Hausdorff distances. Compactness values showed contradictory results; they deteriorated in homogeneous division whereas they remained similar or slightly worse in heterogeneous division.

## 3.3. On the Influence of Disparate Ground Truths

As expected, optimizing separate RatLesNetv2 models with segmentations from different annotators produced more different segmentation masks than when optimizing with segmentations from the same annotator. In other words, the three sets of predictions  $\hat{y}_{1,1}, \hat{y}_{1,2}, \hat{y}_{1,3}$  were similar among themselves in the same manner as  $\hat{y}_{2,1}, \hat{y}_{2,2}, \hat{y}_{2,3}$  (Figure 4B, Annotation 1 and 2), and their differences arise from the stochasticity of ConvNets optimization. In contrast, the shape of the distribution of the Dice coefficients that compare masks derived from RatLesNetv2 models optimized with different annotations (Figure 4B, Mixed) was notably different. Also, Annotation 1 and Mixed Dice coefficients as well as Annotation 2 and Mixed Dice coefficients were significantly different ( $p$ -values  $< 0.002$ ).

In a visual inspection, we observed that Annotation 2 was more approximate, with simpler contours, than Annotation 1. Figure 4A (top row) shows the manual segmentations of the scan with the most disparate annotations and Figure 4A (bottom row) shows the predictions on a scan with the highest Dice coefficient on our baseline study when RatLesNetv2 was trained on the different annotations.

**TABLE 3** | Performance evaluation on the test set after post-processing segregated by lesion stage.

Time-point (scans)	Dice	Compactness	HD
24 h (394)	0.831 ± 0.15	26.539 ± 4.86	3.691 ± 3.53
D3 (135)	0.782 ± 0.12	29.705 ± 8.01	3.067 ± 2.31
D7 (53)	0.790 ± 0.11	40.742 ± 10.65	2.580 ± 2.63
D14 (30)	0.735 ± 0.21	36.018 ± 11.07	3.329 ± 5.06
D21 (40)	0.800 ± 0.11	33.797 ± 6.50	2.546 ± 0.92
D28 (89)	0.593 ± 0.28	29.598 ± 7.46	4.238 ± 5.22
D35 (20)	0.751 ± 0.23	31.203 ± 4.72	4.831 ± 5.76
Shams (107)	1.000 ± 0.00	—	—
2 h (6)	0.719 ± 0.11	23.111 ± 2.27	1.920 ± 0.16
24 h (400)	0.826 ± 0.15	23.218 ± 4.67	3.919 ± 3.79
D3 (129)	0.809 ± 0.10	23.376 ± 5.15	2.796 ± 2.24
D7 (47)	0.860 ± 0.09	23.555 ± 3.99	2.439 ± 2.83
D14 (24)	0.827 ± 0.19	21.705 ± 3.36	3.015 ± 5.69
D21 (34)	0.877 ± 0.10	23.874 ± 2.38	2.002 ± 0.70
D28 (83)	0.692 ± 0.25	22.147 ± 4.65	2.875 ± 1.93
D35 (14)	0.886 ± 0.07	22.037 ± 2.55	1.700 ± 0.67
Shams (131)	1.000 ± 0.00	—	—

Top: Homogeneous division. Bottom: Heterogeneous division.

**TABLE 4** | Ablation study.

Study	Dice (no shams)	Compactness	HD
Baseline	0.784 ± 0.18	29.332 ± 7.86	3.522 ± 3.64
DenseNetBlock <sup>*</sup>	<b>0.771 ± 0.20</b>	<b>30.094 ± 8.86</b>	3.692 ± 3.96
Halving RF	<b>0.754 ± 0.20</b>	30.766 ± 10.57	3.340 ± 3.91
Halving RF <sup>*</sup>	<b>0.765 ± 0.19</b>	<b>31.867 ± 10.41</b>	3.464 ± 3.53
Width-28	0.781 ± 0.18	30.095 ± 8.42	3.423 ± 2.84
Width-36	<b>0.765 ± 0.19</b>	<b>31.620 ± 10.09</b>	3.557 ± 3.73
Baseline	0.813 ± 0.16	23.105 ± 4.58	3.334 ± 3.34
DenseNetBlock <sup>*</sup>	<b>0.801 ± 0.16</b>	<b>23.313 ± 5.13</b>	3.093 ± 2.70
Halving RF	<i>0.819 ± 0.15</i>	<b>25.226 ± 5.34</b>	<b>3.679 ± 3.12</b>
Halving RF <sup>*</sup>	<i>0.820 ± 0.15</i>	<b>25.394 ± 5.51</b>	<b>3.719 ± 3.40</b>
Width-28	<b>0.803 ± 0.17</b>	22.861 ± 4.63	2.892 ± 2.94
Width-36	<b>0.801 ± 0.16</b>	24.036 ± 5.04	2.900 ± 2.84

Top: Homogeneous task. Bottom: Heterogeneous task. Bold: baseline significantly better. Italic: baseline significantly worse ( $p$ -values < 0.05). <sup>\*</sup>Equal number of parameters as Baseline.

### 3.4. The Impact of the Training Set Size on the Performance

**Figure 5** illustrates the evolution of the Dice coefficients, compactness values and Hausdorff distances as the training set increases in size. Dice coefficients (**Figure 5**, left) were remarkably different across time-points and almost every time-point reached a performance plateau with large data sets. Time-points 24 h and D3—which composed the majority of the test set scans by 56.7 and 17.8% of the total, respectively—reached their plateaus later. This effect can be a consequence of the variability within samples. On the contrary, the time-points with the lowest

number of samples (2 h and D35 lesions with 1 and 9 image, respectively) exhibited fluctuations.

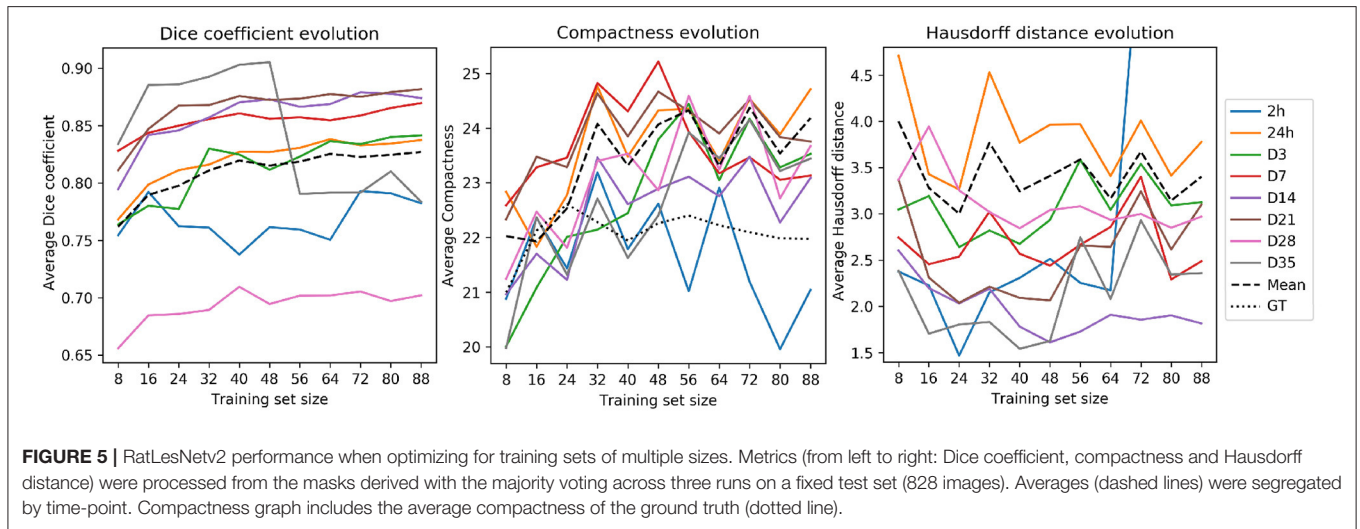
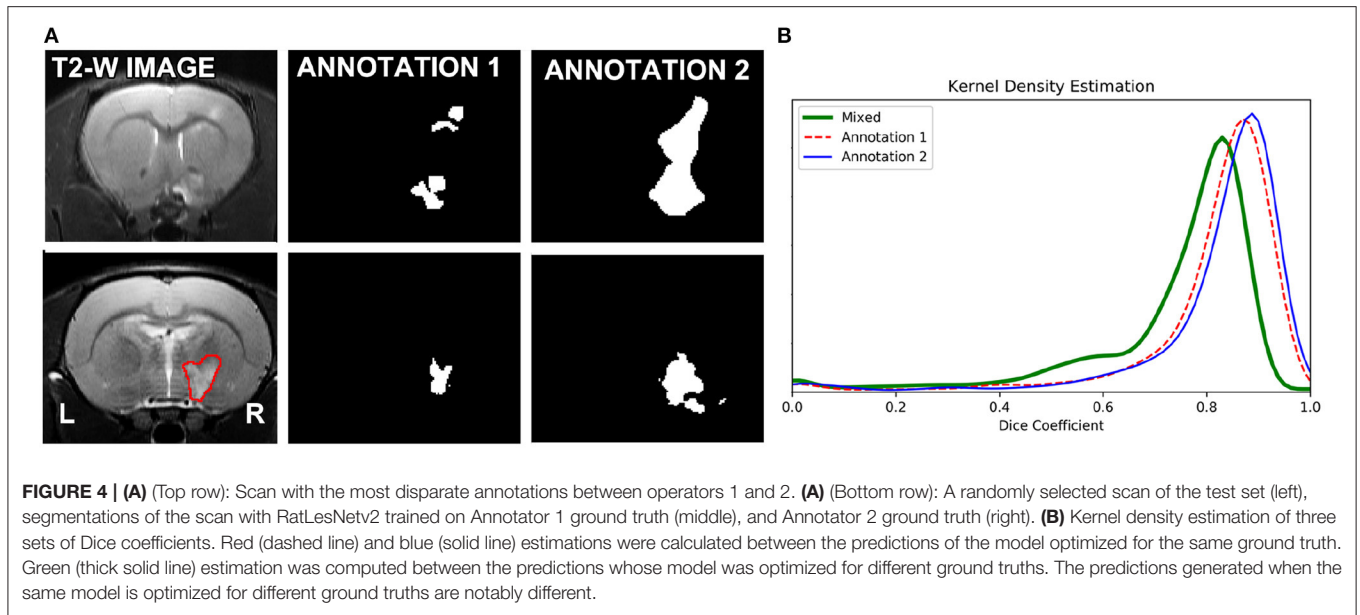
Compactness values (**Figure 5**, center) and Hausdorff distances (**Figure 5**, right) oscillated considerably regardless of the time-point. Hausdorff distances were higher in the time-points with the largest number of samples (24 h and D3), likely due to the existence of outliers. Compactness values, including the average (dashed line), increased analogously to the training set size, i.e., enlarging the training set yielded less compact segmentations. Yet, these compactness values were markedly lower than the compactness values derived from segmentations produced by 3D U-Net, VoxResNet, and RatLesNet (section 3.1).

## 4. DISCUSSION

We showed that RatLesNetv2 yielded similar or better Dice coefficients and Hausdorff distances, and notably more compact segmentations than other convolutional neural networks (Çiçek et al., 2016; Chen et al., 2018a; Valverde et al., 2019). These measurements indicate that the segmentations derived from RatLesNetv2 were more similar to the ground truth, had less large segmentation errors and were more realistic. Additionally, the smaller differences between Hausdorff distances before and after post-processing derived from RatLesNetv2 also indicate that RatLesNetv2 produced fewer segmentation errors far from the lesion surface.

RatLesNetv2 produced more compact segmentations than the other ConvNets without directly minimizing compactness (see **Table 2**), indicating that RatLesNetv2 architecture favors segmentations with smooth borders without holes. Although optimizing compactness (and Hausdorff distance) directly might further improve the results, incorporating these terms to the loss function leads to additional hyper-parameters that require costly tuning. Dice coefficients had large standard deviations and were lower than in existing human brain tumor segmentation studies (Jiang et al., 2019; Myronenko and Hatamizadeh, 2019). These results may arise due to the subjectivity of the segmentation task caused by low image contrast in certain lesions and its consequent high inter- and intra-rater disagreement. However, this is not unexpected as relatively low Dice coefficients and large standard deviations are typical in rodent (Mulder et al., 2017a; Valverde et al., 2019) and human brain lesion segmentation studies (Chen et al., 2017; Valverde et al., 2017; Subbanna et al., 2019), even when studying inter-rater disagreement of manual annotations relying on a semi-automatic segmentation pipeline (Mulder et al., 2017a). We also argued that Dice coefficient alone is not sufficient to measure the segmentation performance. To illustrate the importance of providing additional measurements, consider a brain with a very large and a very small lesion. If the segmentation accurately predicts the large lesion and ignores the small one, Dice coefficients will have a high value not reflecting the segmentation error, but Hausdorff distance is high capturing the segmentation error. Likewise, a lesion segmentation mask with non-smooth surface and several small holes and islands





(i.e., a high compactness value) may have a high Dice coefficient despite being unrealistic.

The difference in the performance between homogeneous and heterogeneous data set divisions indicates that although few 24 h lesion volumes were needed to generalize well, adding more 24 h lesion volumes to the training data (homogeneous division) made RatLesNetv2 specialize on that time-point (Table 3). On the other hand, increasing data diversity (heterogeneous division) improved performance, demonstrating that RatLesNetv2 is capable of learning from a heterogeneous data set. Thus, training on this heterogeneous division increased RatLesNetv2 capability to extrapolate to different-looking ischemic brain lesions. However, without optimizing on additional data, RatLesNetv2 performance on images with other types of lesions, such as tumor lesions, is limited by the lesions' appearance.

The ablation experiments showed that modifications of RatLesNetv2 architecture yielded similar or worse performance,

justifying RatLesNetv2's architectural choices. Despite both residual connections (He et al., 2016) and DenseNetBlocks (Huang et al., 2017) facilitate gradient propagation (Drozdzal et al., 2016), residual connections were preferred over DenseNetBlocks due to their notably higher performance and lower memory requirements. Additionally, a large receptive field empirically demonstrated to increase compactness and reduce large segmentation errors possibly because RatLesNetv2 considers a larger context. The choice of a large receptive field is in agreement with other state-of-the-art ConvNets that achieve large receptive fields by stacking several convolutional layers and/or utilizing dilated convolutions (Chen et al., 2018b).

Our ground-truth disparity experiment confirmed that predictions generated when the same model is optimized for different ground truths are different. Consequently, the quality of the manually-annotated ground truth has a direct impact on the quality of the lesion masks generated automatically. As there

is no unique definition of “lesion,” it may be advantageous for an algorithm to perform differently depending on the labels of the training set. On the other hand, it may also be desirable to design a robust algorithm that performs consistently regardless of some changes in the annotations.

The experiment of training RatLesNetv2 on several training sets of different sizes showed that even with few available training data RatLesNetv2 can generalize well and, despite increasing its performance when optimizing on larger training sets, such improvement is small and compactness values and Hausdorff distances fluctuate considerably.

## 5. CONCLUSION

We presented and made publicly available RatLesNetv2, a 3D ConvNet to segment rodent brain lesions. RatLesNetv2 has been evaluated on an exceptionally large and diverse data set of 916 rat brain MR images, validating RatLesNetv2 reliability on a wide variety of lesion stages with lesions of different appearance. Additionally, RatLesNetv2 produced segmentations that exceeded overall inter-rater agreement Dice coefficients (inter-rater:  $0.73 \pm 0.12$ , RatLesNetv2:  $0.81 \pm 0.16$ ). This enhancement indicates that RatLesNetv2 produces segmentations that are remarkably more consistent with the ground truth than the similarity between different human-made annotations. This consistency is of special importance for research reproducibility, crucial in preclinical studies.

Based on our experiments and, more specifically, the accuracy greater than inter-rater agreement and than of other ConvNets, RatLesNetv2 can be used to automate lesion segmentation in preclinical MRI studies on rats.

## DATA AVAILABILITY STATEMENT

The data analyzed in this study is subject to the following licenses/restrictions: the software and the trained models are publicly available from <https://github.com/jmlipman/RatLesNetv2>. All the quantitative segmentation measures are **Supplementary Material** of this manuscript. The MR image

dataset cannot be shared due to ownership and intellectual property restrictions.

## ETHICS STATEMENT

The animal study was reviewed and approved by National Animal Experiment Board, Finland.

## AUTHOR CONTRIBUTIONS

JV and JT designed the study and analyzed the data. JV conceptualized the algorithm design, developed RatLesNetv2, and wrote the draft. JV, RD, and JT developed the methodology. ASH collected the data. ASH, OG, and ASi interpreted the data. All authors contributed to manuscript revision, proofreading, and approved the submitted version.

## FUNDING

The work of JV was funded from the European Union’s Horizon 2020 Framework Programme [Marie Skłodowska Curie grant agreement #740264 (GENOMMED)] and RD’s work was funded from Marie Skłodowska Curie grant agreement #691110 (MICROBRADAM). We also acknowledge the Academy of Finland grants (#275453 to ASi and #316258 to JT).

## ACKNOWLEDGMENTS

Part of the computational analysis was run on the servers provided by Bioinformatics Center, University of Eastern Finland, Finland. This manuscript has been released as a preprint at arXiv (Valverde et al., 2020).

## SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2020.610239/full#supplementary-material>

## REFERENCES

- Arnaud, A., Forbes, F., Coquery, N., Collomb, N., Lemasson, B., and Barbier, E. L. (2018). Fully automatic lesion localization and characterization: application to brain tumors using multiparametric quantitative mri data. *IEEE Trans. Med. Imaging* 37, 1678–1689. doi: 10.1109/TMI.2018.2794918
- Bribiesca, E. (2008). An easy measure of compactness for 2D and 3D shapes. *Pattern Recogn.* 41, 543–554. doi: 10.1016/j.patcog.2007.06.029
- Chen, H., Dou, Q., Yu, L., Qin, J., and Heng, P.-A. (2018a). Voxresnet: deep voxelwise residual networks for brain segmentation from 3D MR images. *Neuroimage* 170, 446–455. doi: 10.1016/j.neuroimage.2017.04.041
- Chen, L., Bentley, P., and Rueckert, D. (2017). Fully automatic acute ischemic lesion segmentation in DWI using convolutional neural networks. *Neuroimage Clin.* 15, 633–643. doi: 10.1016/j.nicl.2017.06.016
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H. (2018b). “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *Proceedings of the European Conference on Computer Vision (ECCV)* (Munich), 801–818. doi: 10.1007/978-3-030-01234-2\_49
- Choi, C.-H., Yi, K. S., Lee, S.-R., Lee, Y., Jeon, C.-Y., Hwang, J., et al. (2018). 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. *J. Cereb. Blood Flow Metab.* 38, 1371–1383. doi: 10.1177/0271678X17714179
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T., and Ronneberger, O. (2016). “3D U-net: learning dense volumetric segmentation from sparse annotation,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Athens: Springer), 424–432. doi: 10.1007/978-3-319-46723-8\_49
- De Feo, R., and Giove, F. (2019). Towards an efficient segmentation of small rodents brain: a short critical review. *J. Neurosci. Methods* 323, 82–89. doi: 10.1016/j.jneumeth.2019.05.003
- Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology* 26, 297–302. doi: 10.2307/1932409
- Dietterich, T. G. (2000). “Ensemble methods in machine learning,” in *International Workshop on Multiple Classifier Systems* (Cagliari: Springer), 1–15. doi: 10.1007/3-540-45014-9\_1

- Drozdal, M., Vorontsov, E., Chartrand, G., Kadoury, S., and Pal, C. (2016). "The importance of skip connections in biomedical image segmentation," in *Deep Learning and Data Labeling for Medical Applications*, eds G. Carneiro, D. Mateus, P. Loïc, A. Bradley, J. M. R. S. Tavares, V. Belagiannis, J. P. Papa, J. C. Nascimento, M. Loog, Z. Lu, J. S. Cardoso, and J. Cornebise (Athens: Springer), 179–187. doi: 10.1007/978-3-319-46976-8\_19
- Duong, M. T., Rudie, J. D., Wang, J., Xie, L., Mohan, S., Gee, J. C., et al. (2019). Convolutional neural network for automated flair lesion segmentation on clinical brain mr imaging. *Am. J. Neuroradiol.* 40, 1282–1290. doi: 10.3174/ajnr.A6138
- Dutta, S., and Sengupta, P. (2016). Men and mice: relating their ages. *Life Sci.* 152, 244–248. doi: 10.1016/j.lfs.2015.10.025
- Fenster, A., and Chiu, B. (2005). "Evaluation of segmentation algorithms for medical imaging," in *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference* (Shanghai), 7186–7189. doi: 10.1109/IEMBS.2005.1616166
- Gabr, R. E., Coronado, I., Robinson, M., Sujit, S. J., Datta, S., Sun, X., et al. (2019). Brain and lesion segmentation in multiple sclerosis using fully convolutional neural networks: a large-scale study. *Mult. Scler. J.* 26, 1217–1226. doi: 10.1177/1352458519856843
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV), 770–778. doi: 10.1109/CVPR.2016.90
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. (2017). "Densely connected convolutional networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Honolulu, HI), 4700–4708. doi: 10.1109/CVPR.2017.243
- Ioffe, S., and Szegedy, C. (2015). Batch normalization: accelerating deep network training by reducing internal covariate shift. *arXiv* 1502.03167.
- Jiang, Z., Ding, C., Liu, M., and Tao, D. (2019). "Two-stage cascaded U-net: 1st place solution to brats challenge 2019 segmentation task" in *International MICCAI Brainlesion Workshop* (Shenzhen: Springer), 231–241. doi: 10.1007/978-3-030-46640-4\_22
- Kingma, D. P., and Ba, J. (2014). Adam: a method for stochastic optimization. *CoRR* abs/1412.6980.
- Koizumi, J., Yoshida, Y., Nakazawa, T., and 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. *Jpn. J. Stroke* 8, 1–8. doi: 10.3995/jstroke.8.1
- Li, H., Xu, Z., Taylor, G., Studer, C., and Goldstein, T. (2018). "Visualizing the loss landscape of neural nets," in *Advances in Neural Information Processing Systems*, eds S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Montreal, QC), 6389–6399.
- Meyers, R., Lu, M., de Puiseau, C. W., and Meisen, T. (2019). Ablation studies in artificial neural networks. *arXiv* abs/1901.08644.
- Milletari, F., Navab, N., and Ahmadi, S.-A. (2016). "V-net: fully convolutional neural networks for volumetric medical image segmentation," in *2016 Fourth International Conference on 3D Vision (3DV)* (Stanford, CA: IEEE), 565–571. doi: 10.1109/3DV.2016.79
- Moraga, A., Gómez-Vallejo, V., Cuartero, M. I., Szczupak, B., San Sebastián, E., Markuerkiaga, I., et al. (2016). Imaging the role of toll-like receptor 4 on cell proliferation and inflammation after cerebral ischemia by positron emission tomography. *J. Cereb. Blood Flow Metab.* 36, 702–708. doi: 10.1177/0271678X15627657
- Mulder, I. A., Khmelinskii, A., Dzyubachyk, O., de Jong, S., Rieff, N., Wermer, M. J., et al. (2017a). Automated ischemic lesion segmentation in mri mouse brain data after transient middle cerebral artery occlusion. *Front. Neuroinform.* 11:3. doi: 10.3389/fninf.2017.00003
- Mulder, I. A., Khmelinskii, A., Dzyubachyk, O., De Jong, S., Wermer, M. J., Hoehn, M., et al. (2017b). MRI mouse brain data of ischemic lesion after transient middle cerebral artery occlusion. *Front. Neuroinform.* 11:51. doi: 10.3389/fninf.2017.00051
- Myronenko, A., and Hatamizadeh, A. (2019). "Robust semantic segmentation of brain tumor regions from 3D MRIs," in *International MICCAI Brainlesion Workshop* (Shenzhen: Springer), 82–89. doi: 10.1007/978-3-030-46643-5\_8
- Noh, H., Hong, S., and Han, B. (2015). "Learning deconvolution network for semantic segmentation," in *Proceedings of the IEEE International Conference on Computer Vision* (Santiago), 1520–1528. doi: 10.1109/ICCV.2015.178
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). "Pytorch: an imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, eds H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Vancouver, BC), 8024–8035.
- Ronneberger, O., Fischer, P., and Brox, T. (2015). "U-net: convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Munich: Springer), 234–241. doi: 10.1007/978-3-319-24574-4\_28
- Rote, G. (1991). Computing the minimum hausdorff distance between two point sets on a line under translation. *Inform. Process. Lett.* 38, 123–127. doi: 10.1016/0020-0190(91)90233-8
- Roy, S., Knutsen, A., Korotcov, A., Bosomtwi, A., Dardzinski, B., Butman, J. A., et al. (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)* (Washington, DC: IEEE), 687–691. doi: 10.1109/ISBI.2018.8363667
- Santurkar, S., Tsipras, D., Ilyas, A., and Madry, A. (2018). "How does batch normalization help optimization?" in *Advances in Neural Information Processing Systems*, eds S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Montreal, QC), 2483–2493.
- Subbanna, N. K., Rajashekar, D., Cheng, B., Thomalla, G., Fiehler, J., Arbel, T., et al. (2019). Stroke lesion segmentation in flair MRI datasets using customized markov random fields. *Front. Neurol.* 10:541. doi: 10.3389/fneur.2019.00541
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2016). "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV), 2818–2826. doi: 10.1109/CVPR.2016.308
- Valverde, J. M., Shatillo, A., De Feo, R., Gröhn, O., Sierra, A., and Tohka, J. (2019). "Automatic rodent brain mri lesion segmentation with fully convolutional networks," in *International Workshop on Machine Learning in Medical Imaging* (Shenzhen: Springer), 195–202. doi: 10.1007/978-3-030-32692-0\_23
- Valverde, J. M., Shatillo, A., De Feo, R., Gröhn, O., Sierra, A., and Tohka, J. (2020). Ratlesnetv2: a fully convolutional network for rodent brain lesion segmentation. *arXiv* 2001.09138.
- Valverde, S., Cabezas, M., Roura, E., González-Villà, S., Pareto, D., Vilanova, J. C., et al. (2017). Improving automated multiple sclerosis lesion segmentation with a cascaded 3D convolutional neural network approach. *Neuroimage* 155:159–168. doi: 10.1016/j.neuroimage.2017.04.034
- Wang, Y., Cheung, P.-T., Shen, G. X., Bhatia, I., Wu, E. X., Qiu, D., et al. (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. *Int. J. Dev. Neurosci.* 25, 1–5. doi: 10.1016/j.ijdevneu.2006.12.003
- Yang, H., Huang, W., Qi, K., Li, C., Liu, X., Wang, M., et al. (2019). "CLCI-net: cross-level fusion and context inference networks for lesion segmentation of chronic stroke," in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Shenzhen: Springer), 266–274. doi: 10.1007/978-3-030-32248-9\_30

**Conflict of Interest:** As disclosed in the affiliation section, ASH is a full-time payroll employee of the Charles River Discovery Services, Finland—a commercial pre-clinical contract research organization (CRO), which participated in the project and provided raw data as a part of company's R&D initiative.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Valverde, Shatillo, De Feo, Gröhn, Sierra and Tohka. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.