Deep learning-based automated lesion segmentation on mouse stroke magnetic resonance images

Jeehye An[1,2*], Leo Wendt[3*],Georg Wiese[3], Tom Herold[3], Norman Rzepka[3], Susanne Mueller[1,2], Stefan Paul Koch[1,2], Christian J. Hoffmann[1,4], Christoph Harms[1,5-7], Philipp Boehm-Sturm[1,2†]

[1]Charité-Universitätsmedizin Berlin, corporate member of Freie Universität Berlin and Humboldt-Universität zu Berlin, Department of Experimental Neurology and Center for Stroke Research, Charitéplatz 1, 10117 Berlin, Germany.

[2]Charité-Universitätsmedizin Berlin, NeuroCure Cluster of Excellence and Charité Core Facility 7T Experimental MRIs

[3]scalable minds GmbH, Potsdam, Germany

[4]Berlin Institute of Health (BIH), Anna-Louisa-Karsch-Straße 2, 10178, Berlin, Germany

[5]German Center for Cardiovascular Research (DZHK), partner site Berlin

[6]Einstein Center for Neuroscience, Berlin, Germany

[7]NeuroCure Clinical Research Center, Charité-Universitätsmedizin Berlin, Berlin, Germany

*These authors contributed equally to this work

[†] Correspondence:

Philipp Boehm-Sturm

Department of Experimental Neurology, Center for Stroke Research, Charitéplatz 1, Berlin 10117, Germany.

philipp.boehm-sturm@charite.de

Abstract

Magnetic resonance imaging (MRI) is widely used for ischemic stroke lesion detection in mice. A challenge is that lesion segmentation often relies on manual tracing by trained experts, which is labor-intensive, time-consuming, and prone to inter- and intra-rater variability. Here, we present a fully automated ischemic stroke lesion segmentation method for mouse T2-weighted MRI data. As an end-to-end deep learning approach, the automated lesion segmentation requires very little preprocessing and works directly on the raw MRI scans. We randomly split a large dataset of 382 MRI scans into a subset (n=293) to train the automated lesion segmentation and a subset (n=89) to evaluate its performance. We compared Dice coefficients and accuracy of lesion volume against manual segmentation, as well as its performance on an independent dataset from an open repository with different imaging characteristics. The automated lesion segmentation produced segmentation masks with a smooth, compact, and realistic appearance, that are in high agreement with manual segmentation.

Keywords: deep learning, image segmentation, mouse, MRI, stroke

Introduction

In ischemic stroke lesion detection, magnetic resonance imaging (MRI) is frequently used in both clinical and pre-clinical research. T2-weighted (T2w) MR images can be used to detect vasogenic edema and are often used as a proxy of the final infarct since in the hours after ischemic stroke, the blood-brain barrier is disrupted and leads fluid to build up in the surrounding tissue. The resulting extracellular water increases T2 relaxation time, leading to increased signal in the stroke area. Ischemic brain tissue damage can be observed from T2w MR images as early as 3.5 hours after the onset of stroke.[1] Edema then peaks at around 24-48h after stroke[2], and stroke lesion volume measurement on T2w MR images in this phase has become an important outcome measure in rodent stroke research.

A challenge in stroke MRI research is that lesion segmentation often relies on manual tracing by trained experts. Manual segmentation is labor-intensive, time-consuming, and prone to inter- and intra-rater variability.[3–5] This becomes especially problematic in pre-clinical research, which may involve large numbers of MR scans or multi-center datasets. Automated lesion segmentation is therefore crucial for both efficiency and reliability. In addition, automated lesion segmentation is essential for developing a fully automated processing of rodent stroke MR images. Atlas registration in the presence of a stroke lesion is challenging, and we have previously shown that this can only be effectively done if an algorithm is informed with a lesion mask.[6]

Several automated lesion segmentation algorithms have been developed[7,8] for human data. With regards to animal data, there are only a few (semi-) automated methods. Jacobs et al.[9] developed an unsupervised lesion segmentation model with an iterative self-organizing data analysis technique (ISODATA), a K-means derived clustering technique. A multiparametric ISODATA was applied to MR data to characterize ischemic lesion tissue in rats. Ghosh et al.[10,11] developed an automated ischemic lesion segmentation algorithm for neonatal rat brains using a hierarchical recursive region splitting (HRS) approach. Mulder et al.[3] developed a level-set-based lesion segmentation algorithm that requires one T2w MR sequence as input. Castaneda-Vega et al.[12] trained a random forest classifier for ischemic stroke lesion segmentation. However, most of these methods require some degree of manual intervention, require different parameters for each dataset for optimal performance, and are often dependent on the performance of other steps, such as preprocessing. Valverde et al.[5] presented RatLesNetv2, a Convolutional neural network (CNN)-based Rodent Brain Lesion Segmentation algorithm for T2w MR images. However, RatLesNetv2 requires users to train their own model using their data.

Automated Mouse Stroke Lesion Segmentation

In this exploratory study, we developed a fully automated ischemic stroke lesion segmentation method for mouse T2w MR images. The automated lesion segmentation builds on top of U-Net[13], a convolutional neural network that is commonly used in biomedical image segmentation. We trained and evaluated the automated segmentation method on a large dataset of 382 T2w mouse MR images. Our approach produced realistic segmentations in high agreement with manually segmented lesion masks. The dataset used in this study is openly available on zenodo (https://doi.org/10.5281/zenodo.6379878) and presents one of the largest public T2w mouse stroke lesion repositories to date. Thus, it will help to make preclinical stroke lesion segmentation on MR images more standardized across studies and labs.

Materials and methods

Animals

We retrospectively reanalyzed male C57/BL6 J mice (Janvier, Germany) from previous stroke studies, most of which are unpublished. Animals underwent middle cerebral artery occlusion (MCAo, n = 332), or a sham procedure (n = 50) at the age of 10–12 weeks. Experimenters were blinded to the condition of the animals. Mice were housed in a temperature (22 ± 2°C), humidity (55 ± 10%), and light (12/12-h light/dark cycle) controlled environment. All surgeries were performed under 1.5%–2% isoflurane in a 70% nitrous oxide and 30% oxygen mixture. During surgery, core body temperature was maintained at 37 ± 0.2°C with an automated rectal probe and heat blanket. After surgery, a topical application of 1% bupivaccain gel was applied on surgical wounds. All experiments were approved by the Landesamt für Gesundheit und Soziales (registration numbers G0197/12, G0005/16, G0057/16, G0119/16, G00254/16, G0157/17, and G0343/17) and were performed in accordance with the German Animal Welfare Act and the ARRIVE.[14]

MCAo

After anesthesia induction, a midline neck incision was made to expose the left carotid artery. A filament (190 µm diameter, Doccol, Sharon, MA/USA) was introduced into the external carotid artery (ECA) and advanced to the origin of the left MCA via the internal carotid artery. After 45 min, the filament was withdrawn. For sham animals, the filament was advanced to the MCA and withdrawn immediately. Both the common carotid artery and the ECA were ligated and remained occluded during reperfusion.[15] All MCAo were performed by two, experienced surgeons.

MRI measurements

Mice were anesthetized with isoflurane (2-2.5% at initiation and 1-2% during the scanning period) in a 70:30 nitrous oxide:oxygen mixture. Respiration rate was monitored with a pressure-sensitive pad using MRI compatible equipment (Small Animal Instruments, Inc., Stony Brook, NY, USA). T2w images were acquired 24 h post stroke surgery on a 7 T MR scanner (Bruker, Ettlingen, Germany) and a 20 mm inner diameter transmit/receive quadrature volume coil (RAPID Biomedical, Rimpar, Germany). A 2D RARE sequence was used (repetition time/echo spacing/effective echo time = 4.2 s/12 ms/36 ms, RARE factor = 8, 32 consecutive 0.5 mm thick

slices, FOV = (25.6 mm^2), image matrix = 256 × 196 zerofilled to 256 × 256, 4 averages, scan time 6:43 min).

Manual lesion segmentation

The manual lesion masks were created in a semi-automated manner by tracing hyperintense lesions using a thresholding function in ANALYZE 10.0 software (AnalyzeDirect, Overland Park, KS, USA). The lesion masks were of the same dimensionality as the T2w images and were exported in NIfTI format. The manual segmentation was performed by an experienced researcher with over 20 years of experience and whose tracings were validated against histology in previous studies[6]. Each manual segmentation took around 5 minutes.

Independent dataset

An independent dataset from Mulder et al.[16] was used to further evaluate the automated lesion segmentation. The independent dataset is different from that of the present study in that it is of lower resolution (128x128 and 196x196, instead of 256x256), it contains quantitative T2 maps (instead of T2w images), and hads different scan times after the infarct induction. The authors provide two manual annotations from independent observers, which were segmented by applying various thresholding operations in ImageJ in a semi-automated manner. Some of the low signal-to-noise ratio data in the dataset therefore contained very fragmented lesion masks with small holes (order of one or few voxels). We evaluated the automated lesion segmentation on a subset of 19 higher quality datasets (Cologne-Set 1, Cologne-Set 2), which we found to have lesion masks that were less fragmented, and thus more realistic, than other datasets.

Automated lesion segmentation

Network architecture

In order to automatically segment the lesion, a variant of the popular U-Net was trained.[13] The architecture follows an encoder-decoder design but adds "skip connections" between intermediate layers of the same spatial resolution.

As shown in Figure 1, two max-pooling layers were applied, allowing the network to analyze the image at three distinct spatial resolutions. The upsampling was implemented as nearest-neighbor upsampling, followed by a convolutional layer with kernel size 1x1x1 and ReLU activation. Between consecutive downsampling or upsampling steps, two convolutions with kernel size 3x3x1 and ReLU activation were applied. At last, a 3x3x3 convolution with ReLU activation was applied, followed by a 1x1x1 convolution with a softmax activation, the output of which is interpreted as the background and lesion probabilities. All in all, the network maps a 196x152x30 image tensor to a 154x110x28x2 tensor of probabilities. The receptive field of the network is 43 x 43 x 3 voxels which equals 4.3 x 4.3 x 1.5 mm^3.

The number of feature maps was set to 8, 16, and 32, depending on the level of the U-Net. The number of feature maps of the next level was already used on the last convolution before the down-sampling.[17] Batch normalization layers were purposefully excluded because we experienced issues in low-contrast parts of the dataset when batch normalization was included. The network had

42,442 trainable parameters.

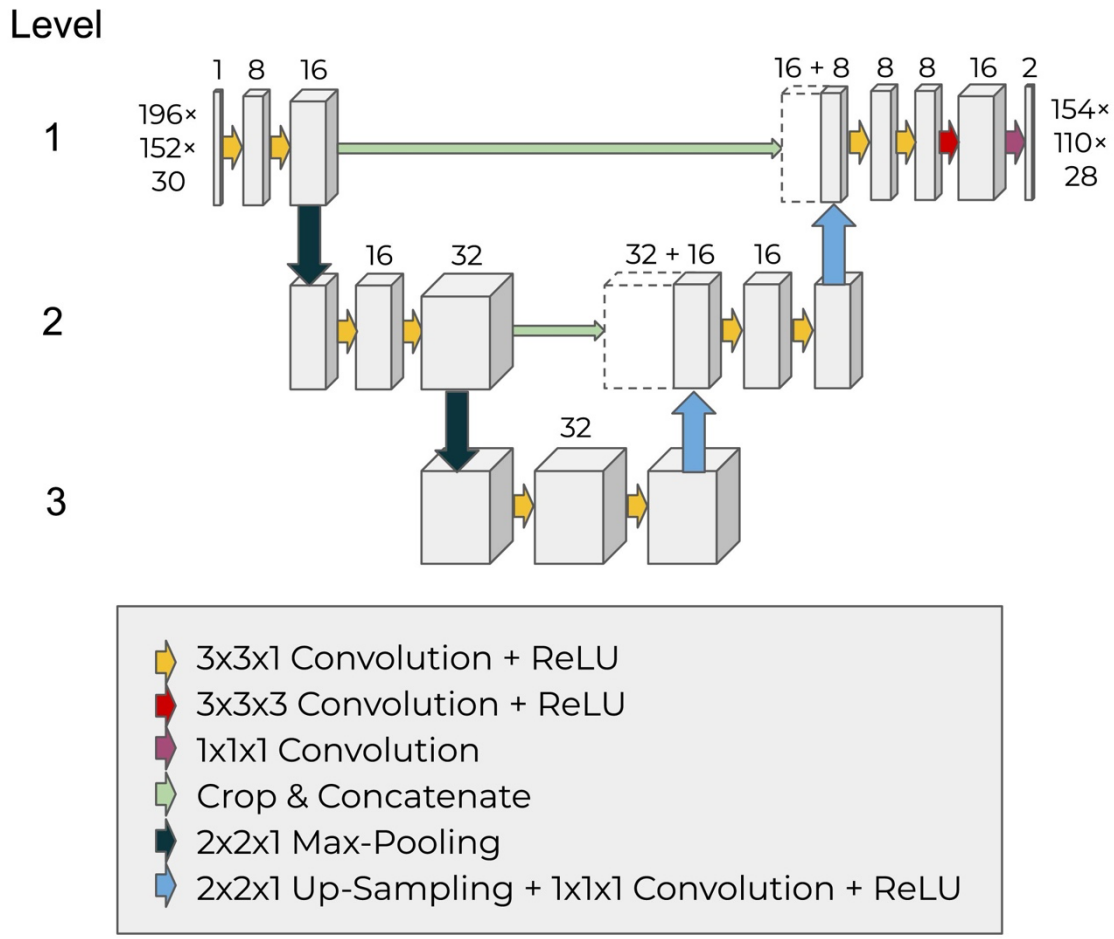

Figure 1. Network Architecture and processing pipeline of automated method.

Model training

The boxes were split randomly into a training (80%) and validation set (20%). The (0, 255) interval of intensity values was scaled to the (-1, 1) interval by transforming them via I := (I - 127.5) / 127. As a data augmentation, voxel-wise Gaussian noise was added to the training samples with a mean of 0 and a standard deviation of 0.45.

The voxel-wise loss was computed as the categorical cross-entropy between the predicted probabilities and the one-hot encoded ground truth class. The training loss was defined as the total categorical cross-entropy across all voxels in a batch of eight training examples. Voxels were weighted inversely proportional to the frequency of the respective class.

All kernels were initialized according to Glorot & Bengio.[18] Biases were initially set to zero. The network was trained to minimize the training loss using the Adam optimizer[19] with a learning rate of 1e-4 and a batch size of 8.

Implementation

The automated lesion segmentation method was implemented in TensorFlow and Python using the Voxelytics framework (https://voxelytics.com; scalable minds, Potsdam, Germany). Training and predictions were executed on a system with an NVidia Geforce GTX1080, an Intel Core i7-6700 CPU @ 3.40GHz, and 64GB RAM.

Qualitative comparison

T2w images and corresponding lesion masks were registered to Allen mouse brain atlas using ANTx2 as described previously (https://github.com/ChariteExpMri/antx2).[6] Incidence maps were generated by displaying voxel-wise mean of lesion masks across animals (expressed in percent) in atlas space. For each animal, a map of false positives was generated by voxel-wise subtraction of the manually generated lesion mask from automatically generated lesion mask in atlas space followed by removal of all values <1. Incidence maps of false positives were then generated by voxel-wise mean of false positive maps across animals (expressed in percent). For incidence maps of false negatives the same procedure was performed except the automatically generated mask was subtracted from the manually generated mask.

Quantitative comparison and statistical analysis

We evaluated the automated lesion segmentation on a dedicated evaluation set consisting of 20% of the overall dataset (89 scans). Two metrics were used for assessment: total lesion volume and Dice coefficient. Total stroke volume is one of the most important readout parameters in preclinical stroke studies and was compared between methods using Pearson correlation and Bland-Altman analyses. Dice coefficient[20] measures the overlap volume between two regions. It is popularly used for performance evaluation in the field of image segmentation.

Results

The resulting segmentation masks accurately detected lesions and produced a smooth, human annotator-like appearance. Incidence maps of manually segmented and automated lesion masks are shown in Figure 2A. We also calculated false positives and false negatives across all evaluation datasets by subtracting manual lesion masks from automated lesion masks, shown in Figure 2B. Qualitative inspection of the spatial distribution showed elevated incidence of false positives at the borders between cortical and subcortical structures and more homogenous distribution of false negatives across the lesion territory (Figure 2B).

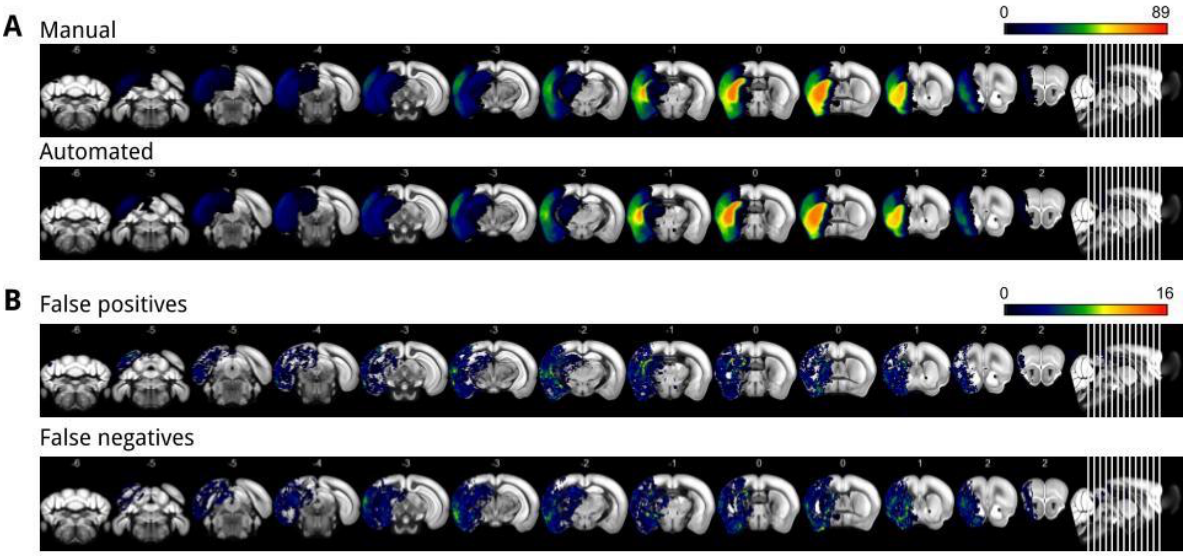

Figure 2. Incidence maps of (A) manual and automated lesion masks, and (B) false positives and false negatives across all evaluation datasets (n = 89).

The evaluation set included 12 scans without any lesions. Of those, the automated lesion segmentation detected small lesion masks in 10 datasets, with an average lesion volume of 1.00 mm^3.

Figure 3A shows the Pearson correlation between manual and automated lesion mask volumes. The plot shows a very high level of correlation, with R = .99. The corresponding Bland-Altman plot (Figure 3B) indicates very high agreement between the two methods. No large systematic error was

found but the automated algorithm tended to underestimate lesion volume for few of the smaller lesions. Overall, we report an average Dice score of 0.91.

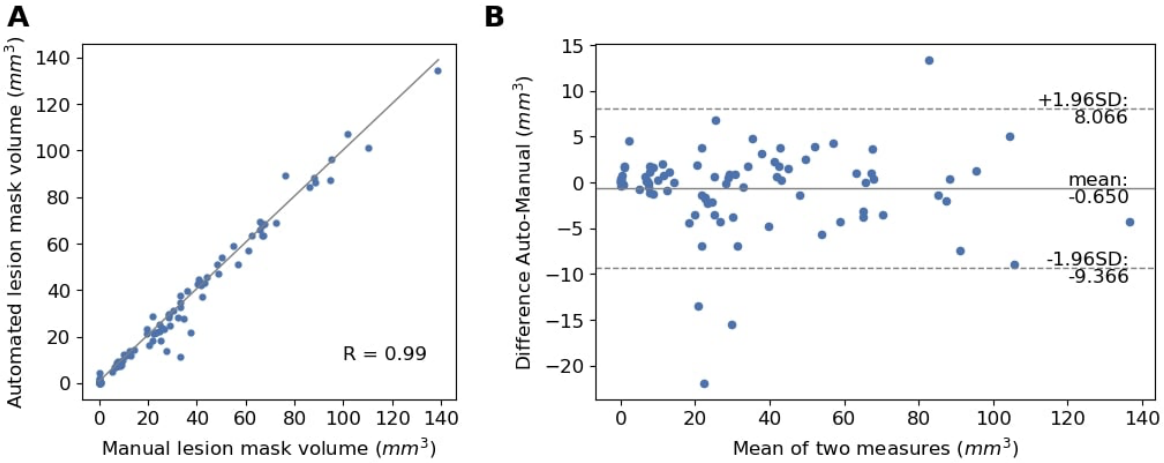

Figure 3. (A) Correlation of manual vs. automated lesion mask volumes. (B) Bland-Altman plot comparing agreement between the manual and automated lesion mask volumes.

To test the generalizability of our automated lesion segmentation method on other datasets, we evaluated our approach on an independent dataset[16] and compared it with the performance of their method. The results are shown in Table 1. Overall, the automated segmentation was able to segment the lesions reliably and with high accuracy. Our lesion masks tended to be smaller and more compact than lesion masks from their method, as shown in Figure 4. We had an average Dice score of 0.76 compared to 0.86.

Age (months)	Infarct induction	Mulder et al. method		Present study	
		Observer 1	Observer 2	Observer 1	Observer 2
3	18h (n=6)	0.88 ± 0.03	0.89 ± 0.03	0.79 ± 0.05	0.81 ± 0.05
	4d (n=3)	0.85 ± 0.01	0.87 ± 0.01	0.68 ± 0.04	0.72 ± 0.03
2, 12	48h (n=10)	0.86 ± 0.07	0.85 ± 0.07	0.82 ± 0.06	0.76 ± 0.08
Average		0.86	0.87	0.76	0.76

Table 1. Comparison of performances of automated lesion segmentation algorithm by Mulder et al.[16] vs. algorithm of the present study on open data repository of Mulder et al.[16] Observer 1 vs. observer 2 refer to lesion masks produced by each of the two manual tracers from their study. Mean dice coefficient across animals within each group and average of mean dice coefficient across three groups are shown. For the present study, errors refer to standard deviation. For Mulder et al.,[3] the type of error was not reported.

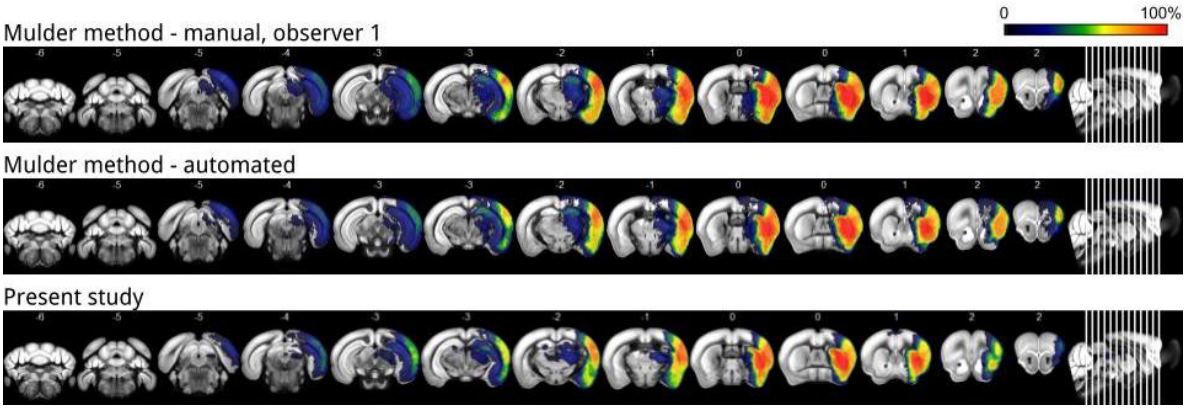

Figure 4. Incidence maps of automated lesion masks generated by Mulder et al.[3] and by the present study.

Discussion

We have presented an automated algorithm for ischemic lesion segmentation from mouse T2w MR images. Our reported average Dice score of 0.91 shows that the automated lesion segmentation produces realistic segmentations that very closely resemble human annotations. Previous studies showed that annotations made by two human annotators can have a Dice score as low as 0.73 or 0.79.[3,5]

The convolutional neural network architecture used in our approach is based on a variation of the 3D U-Net[13] design, a highly successful segmentation model for bioimaging tasks. As an end-to-end deep learning approach, the automated lesion segmentation requires very little preprocessing and works directly on the raw MRI scans. Unlike some of the previous automated lesion segmentation approaches[3,10], the present study is fully automated and it is not dependent on other preprocessing steps, such as registration or skull stripping. This is significant for achieving a fully automated pipeline for processing rodent stroke MR images, especially since we have previously shown that such preprocessing steps are prone to error in presence of a stroke lesion.[6]

Our approach is similar in design to the recently published RatLesNetv2-architecture.[5] However, we used a much larger training set and incorporated 3D information across the z-axis into our predictions. They report dice scores (inter-rater: 0.73 ± 0.12, RatLesNetv2: 0.81 ± 0.16) on their evaluation set that are lower than on our best performing evaluation score (0.91).

We further tested the automated lesion segmentation on the dataset provided by Mulder, Khmelinskii, Dzyubachyk, et al.[16] and tested generalizability across scans with different imaging characteristics (lower resolution; T2 maps instead of T2w images; different scan times after infarct induction). Our Dice score on Mulder, Khmelinskii, Dzyubachyk, et al.[16]'s data fell short of their performance (Table 1). However, their ground-truth annotations contain holes, rough and fringed edges, and spots of isolated voxels since a fixed intensity threshold based on contralateral mean and standard deviation was used without further postprocessing. These differ considerably from ground-truth labels in our dataset, which mostly contain compact lesion masks with smooth edges. We aimed for compact lesion masks based on two arguments. First, consistent with literature[5,21], we argue compact lesion masks are more realistic and comparable to manual annotations since hypoperfusion in large vessel stroke varies smoothly in space and not at the order of single voxels. Second, our manual lesion masks were obtained using a semi-automated manner using a

thresholding function and the lesion masks obtained through the same procedure have been validated by comparison to histological data in a previous study[6].

Another advantage of our automated method over the approach by Mulder et al is that while the latter only takes T2 maps as input, the current method can take any image in NIfTI format as input on which the stroke appears hyperintense, including T2w images, T2 maps, and likely also diffusion-weighted images (DWI).

T2w images on their own may not be the gold standard for estimating final stroke lesion volume since swelling due to post-stroke edema leads to a significant overestimation of lesion volume in T2w images compared to histology[12,22]. Lesion size thus needs to be corrected for the effect of brain edema to avoid an overestimation of the lesion volume. This correction could be done using the atlas registration approach by Koch et al.[6] or the model by Gerriets et al.[22]

In our study, we used T2w images for lesion segmentation. For ischemic stroke segmentation, Diffusion Weighted Imaging (DWI), in addition to T2w images, is also routinely applied. Whereas T2w images are sensitive to vasogenic edema, apparent diffusion coefficient (ADC) maps derived from DWI indicate cytotoxic edema, i.e. water diffusion restrictions due to cellular swelling. ADC and T2w data measure partly overlapping and distinct microstructural lesion traits. It is thus suggested that a co-examination of ADC and T2w data provides a more accurate predictor of final stroke volume.

Despite these limitations, T2w images are still frequently used in preclinical stroke research and are the most commonly used imaging technique to determine the final infarct in animal studies of experimental stroke[23].

To conclude, a fully automated end-to-end deep learning approach to segment stroke lesions in preclinical studies has been trained on one of the largest datasets of T2w MR images acquired in subacute murine stroke to date. The method performed well compared to manual segmentation and will help to standardize MRI-based lesion volumetry between studies and labs. It can now be used to inform other algorithms needed for rodent MR image preprocessing that previously relied on manual interaction.

Acknowledgments

Excellent technical assistance was provided by Janet Lips, Larissa Mosch, Monika Dopatka and Marco Foddis. Ulrich Dirnagl initiated the collaboration between Charité-Universitätsmedizin Berlin and scalable minds GmbH.

Author Contributions

JA, TH, NR, PBS. designed the study, CH, CJH and SM provided and annotated data, JA, TH, NR, SPK, PBS preprocessed and organized the data, TH, NR, GW, LW wrote the software and analyzed data, JA and SPK performed the statistical analyses, JA, GW, TH and PBS wrote the manuscript, all authors commented to the manuscript at all stages.

Conflict of Interest

The automated lesion segmentation algorithm is a product of scalable minds GmbH. TH, LW, GW and NR are employed by scalableminds GmbH. All other 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.

Funding

This study was supported by the Brandenburg Ministry of Economic Affairs, Labour and Energy (Brandenburgischer Innovationsgutschein) and the European Regional Development Fund. Funding to JA, SM, SPK and PBS was provided by the German Federal Ministry of Education and Research under the ERA-NET NEURON scheme (BMBF, 01EW1811), the German Research Foundation (DFG, Project HO5277/3-1, project number 417284923 to CJH, HA5741/5-1, project number 417284923 to CH, Project-ID 424778381 –TRR 295 to CH, BO4484/2-1, project number 428869206 to PBS, the Open Access Publication Fund of Charité - Universitätsmedizin Berlin and EXC NeuroCure) and by a grant from the Fondation Leducq (17 CVD 03) to CH. Retrospective analyses of previously acquired animal data were supported by Charité 3^{R}| Replace - Reduce – Refine.

Data Availability Statement

The datasets used for this study can be found in a zenodo repository (https://doi.org/10.5281/zenodo.6379878).

References

1. Loubinoux I, Volk A, Borredon J, et al. Spreading of vasogenic edema and cytotoxic edema assessed by quantitative diffusion and T2 magnetic resonance imaging. *Stroke* 1997; 28: 419–427.

2. Slivka A, Murphy E, Horrocks L. Cerebral edema after temporary and permanent middle cerebral artery occlusion in the rat. *Stroke* 1995; 26: 1061–1065.

3. Mulder IA, Khmelinskii A, Dzyubachyk O, et al. Automated ischemic lesion segmentation in MRI mouse brain data after transient middle cerebral artery occlusion. *Frontiers in Neuroinformatics* 2017; 11: 3.

4. Rekik I, Allassonnière S, Carpenter TK, et al. Medical image analysis methods in MR/CT-imaged acute-subacute ischemic stroke lesion: Segmentation, prediction and insights into dynamic evolution simulation models. A critical appraisal. *NeuroImage: Clinical* 2012; 1: 164–178.

5. Valverde JM, Shatillo A, de Feo R, et al. RatLesNetv2: A Fully Convolutional Network for Rodent Brain Lesion Segmentation. *Frontiers in Neuroscience* 2020; 0: 1333.

6. Koch S, Mueller S, Foddis M, et al. Atlas registration for edema-corrected MRI lesion volume in mouse stroke models. *Journal of Cerebral Blood Flow and Metabolism* 2019; 39: 313–323.

7. Ito KL, Kim H, Liew S. A comparison of automated lesion segmentation approaches for chronic stroke T1-weighted MRI data. *Human Brain Mapping* 2019; 40: 4669.

8. Lee E-J, Kim Y-H, Kim N, et al. Deep into the Brain: Artificial Intelligence in Stroke Imaging. *Journal of Stroke* 2017; 19: 277.

9. Jacobs MA, Zhang ZG, Knight RA, et al. A model for multiparametric MRI tissue characterization in experimental cerebral ischemia with histological validation in rat: Part 1. *Stroke* 2001; 32: 943–949.

10. Ghosh N, Recker R, Shah A, et al. Automated ischemic lesion detection in a neonatal model of hypoxic ischemic injury. *Journal of Magnetic Resonance Imaging* 2011; 33: 772–781.

11. Ghosh N, Sun Y, Bhanu B, et al. Automated detection of brain abnormalities in neonatal hypoxia ischemic injury from MR images. *Medical Image Analysis* 2014; 18: 1059–1069.

12. Castaneda-Vega S, Katiyar P, Russo F, et al. Machine learning identifies stroke features between species. *Theranostics* 2021; 11: 3017.

13. Ronneberger O, Fischer P, Brox T. U-Net: Convolutional Networks for Biomedical Image Segmentation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 2015; 9351: 234–241.

14. Kilkenny C, Browne WJ, Cuthill IC, et al. Improving bioscience research reporting: The arrive guidelines for reporting animal research. *PLoS Biology*; 8. Epub ahead of print June 2010. DOI: 10.1371/journal.pbio.1000412.

15. O E, S K, U D, et al. Modeling stroke in mice - middle cerebral artery occlusion with the filament model. *J Vis Exp*. Epub ahead of print 2011. DOI: 10.3791/2423.

16. Mulder IA, Khmelinskii A, Dzyubachyk O, et al. MRI Mouse Brain Data of Ischemic Lesion after Transient Middle Cerebral Artery Occlusion. *Frontiers in Neuroinformatics* 2017; 0: 51.

17. Çiçek Ö, Abdulkadir A, Lienkamp SS, et al. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 2016; 9901 LNCS: 424–432.

18. Glorot X, Bengio Y. Understanding the difficulty of training deep feedforward neural networks. 2010; 249–256.

19. Kingma DP, Ba J. Adam: A Method for Stochastic Optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, https://arxiv.org/abs/1412.6980v9 (2014, accessed October 5, 2021).

20. Dice LR. Measures of the Amount of Ecologic Association Between Species. *Ecology* 1945; 26: 297–302.

21. Fenster A, Chiu B. Evaluation of segmentation algorithms for medical imaging. *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings* 2005; 7 VOLS: 7186–7189.

22. Gerriets T, Stolz E, Walberer M, et al. Noninvasive Quantification of Brain Edema and the Space-Occupying Effect in Rat Stroke Models Using Magnetic Resonance Imaging. *Stroke* 2004; 35: 566–571.

23. Milidonis X, Marshall I, Macleod MR, et al. Magnetic resonance imaging in experimental stroke and comparison with histology: Systematic review and meta-analysis. *Stroke* 2015; 46: 843–851.

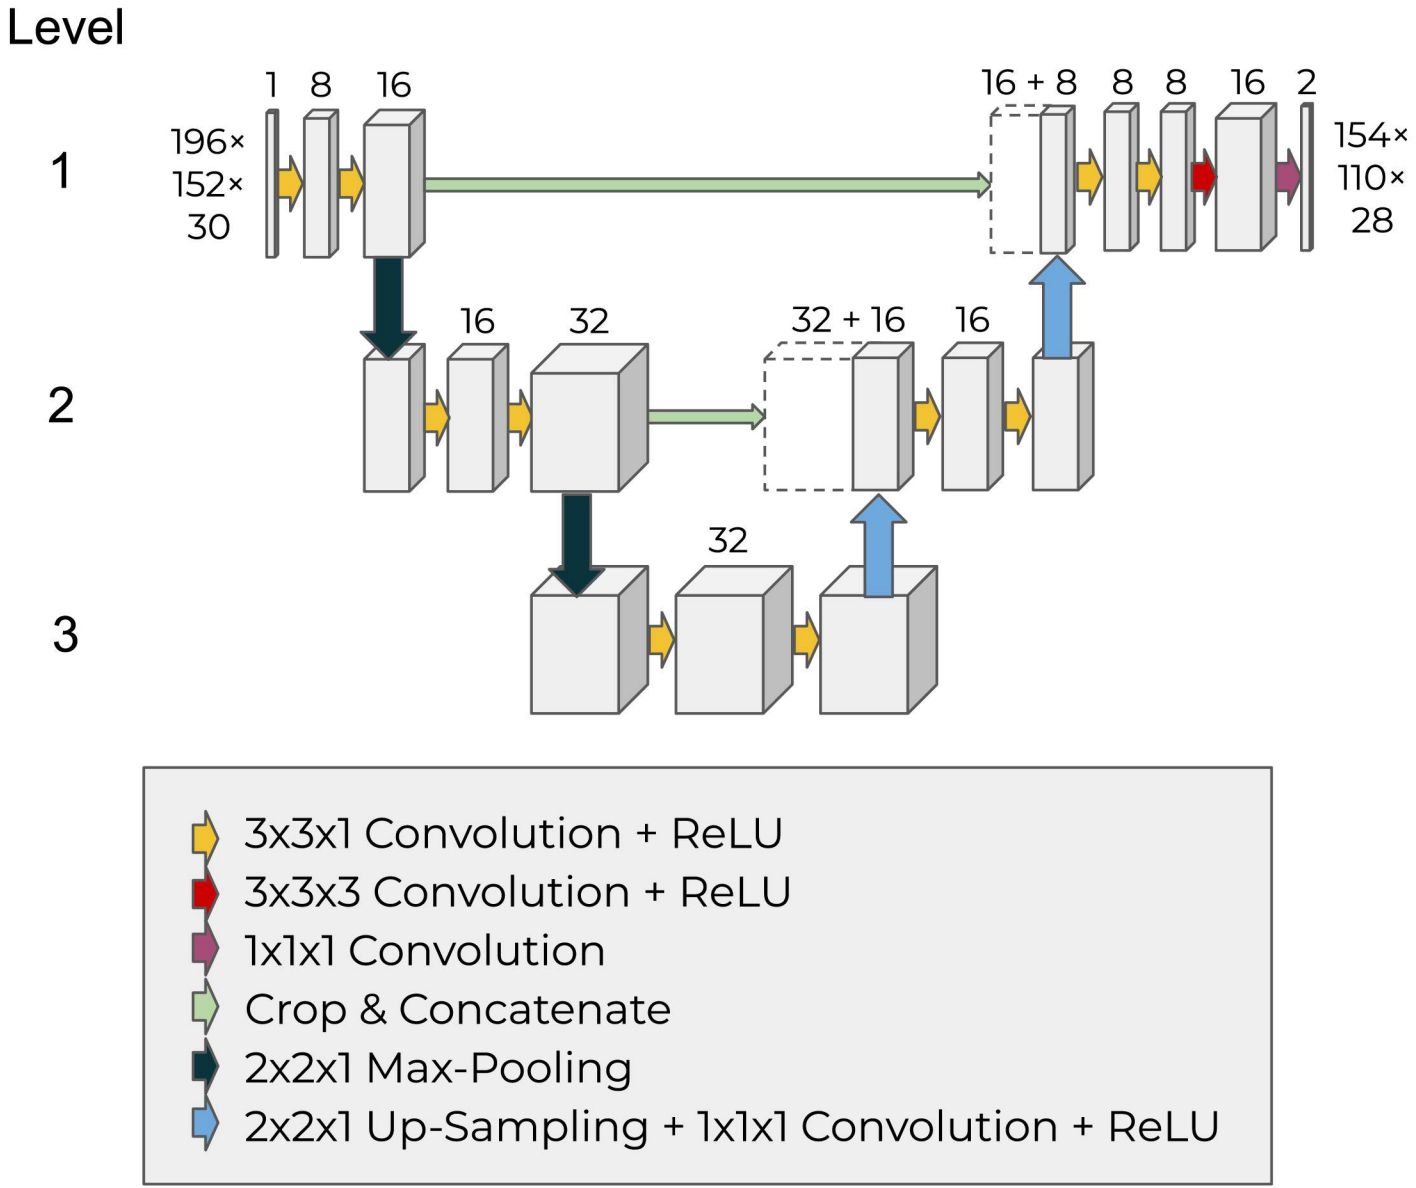
Level
1
2
3
196×
152×
30
1
8
16
16
32
32
16 + 8
8
8
16
2
154×
110×
28
32 + 16
16
3x3x1 Convolution + ReLU
3x3x3 Convolution + ReLU
1x1x1 Convolution
Crop & Concatenate
2x2x1 Max-Pooling
2x2x1 Up-Sampling + 1x1x1 Convolution + ReLU

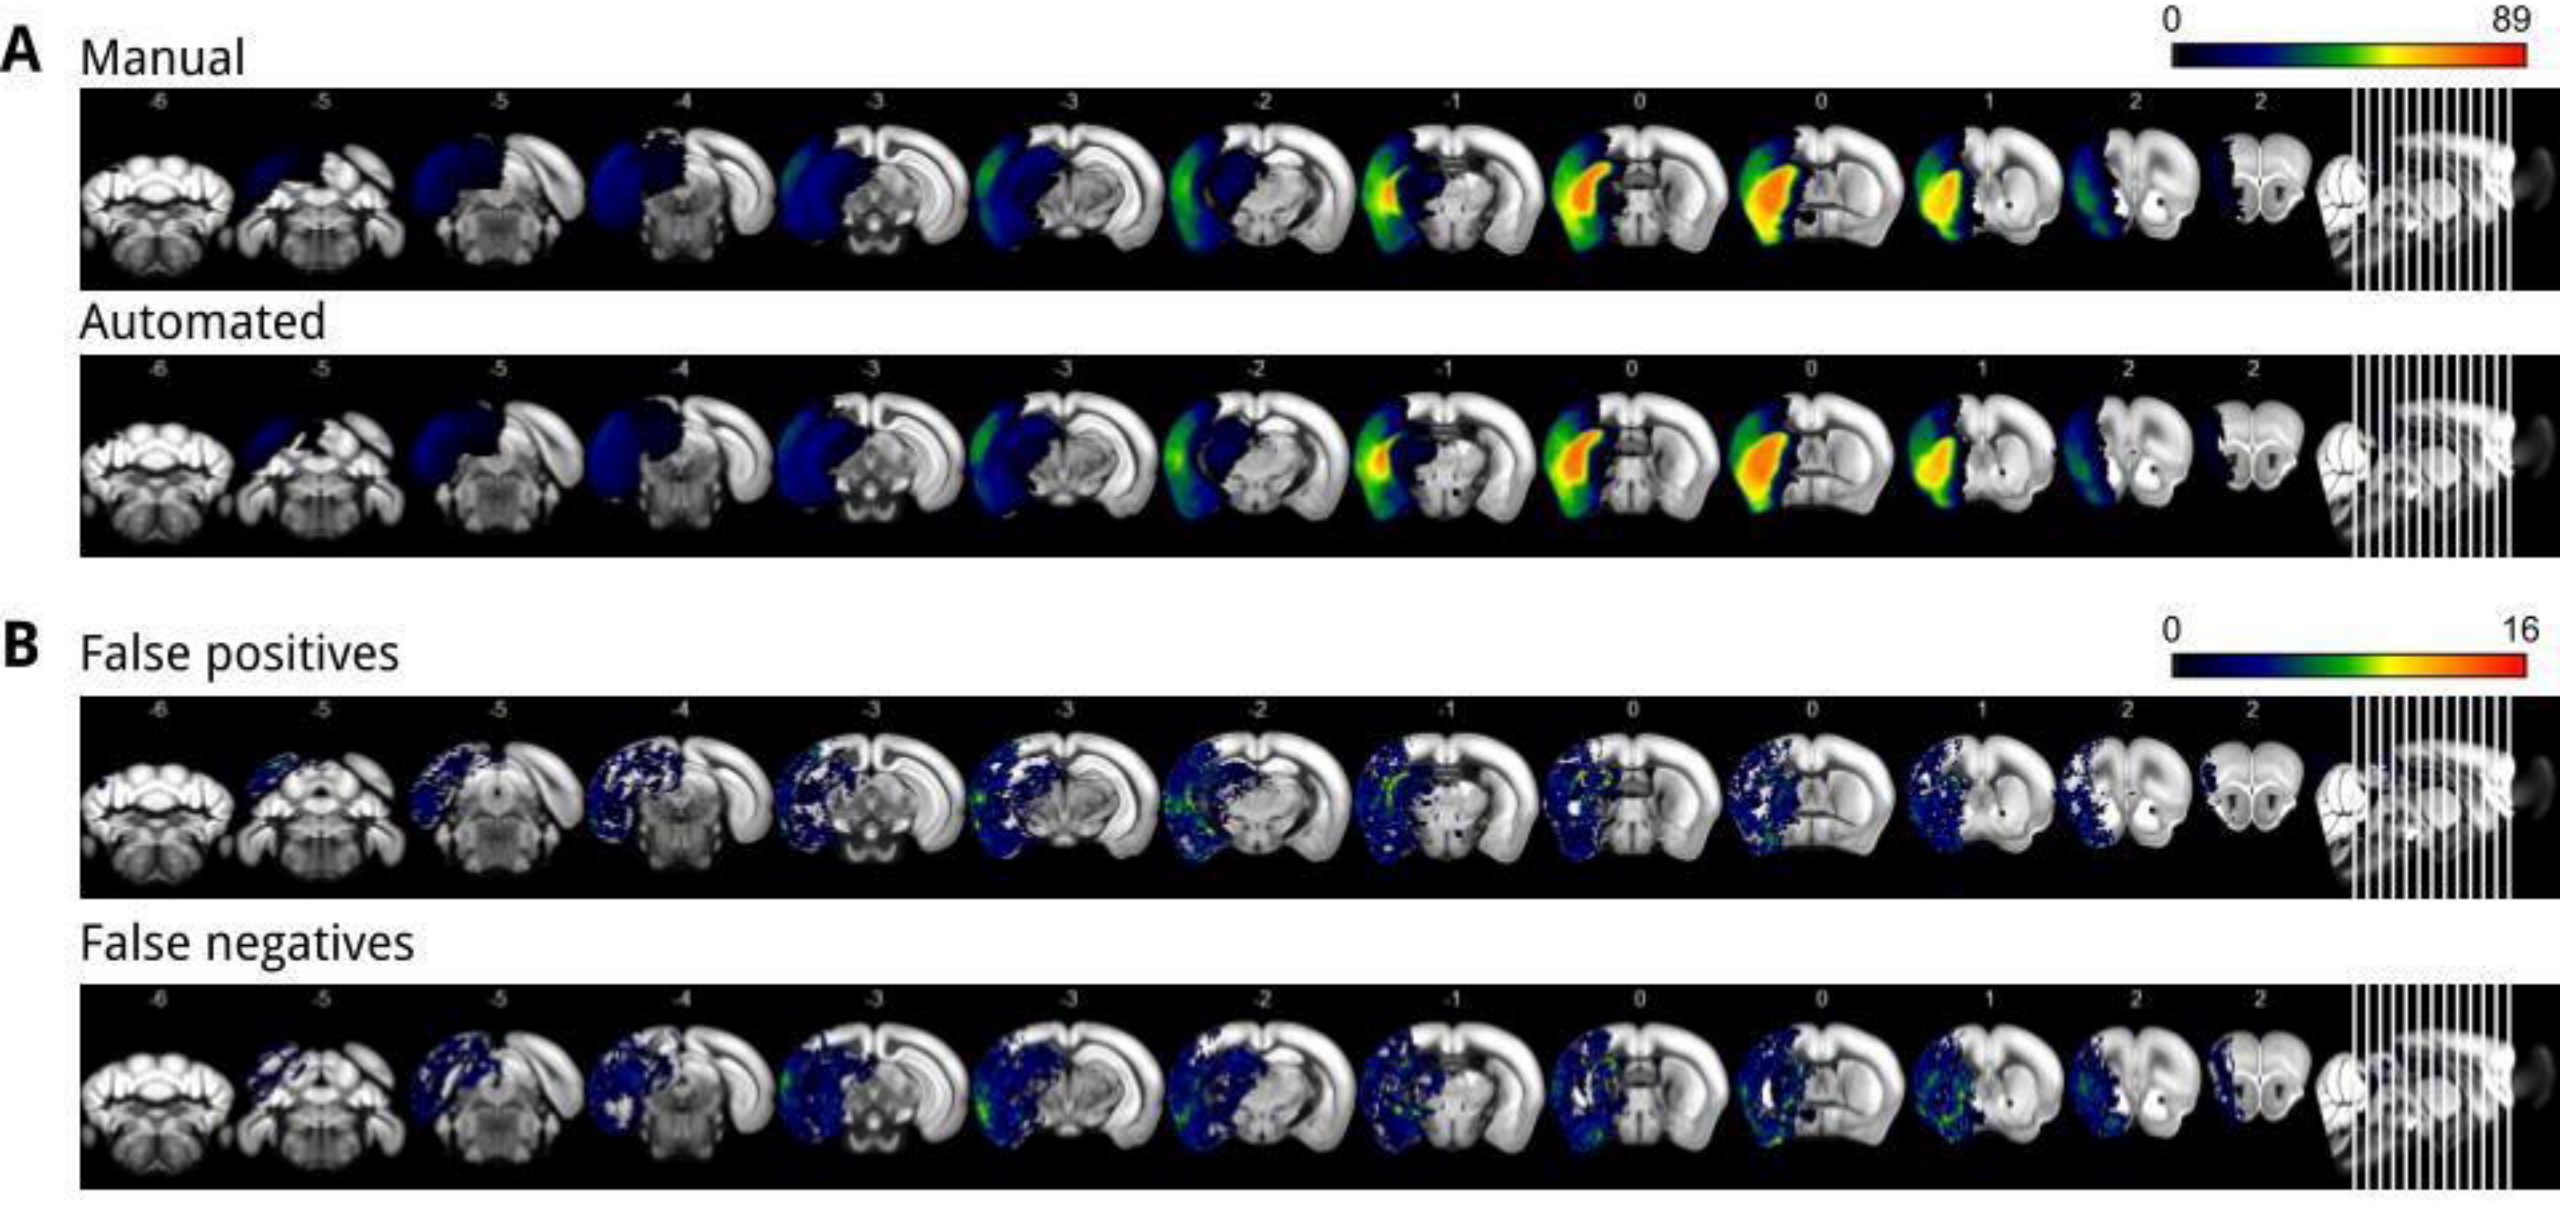
A
Manual
0
89
-6
-5
-5
-4
-3
-3
-2
-1
0
0
1
2
2
Automated
B
False positives
0
16
False negatives

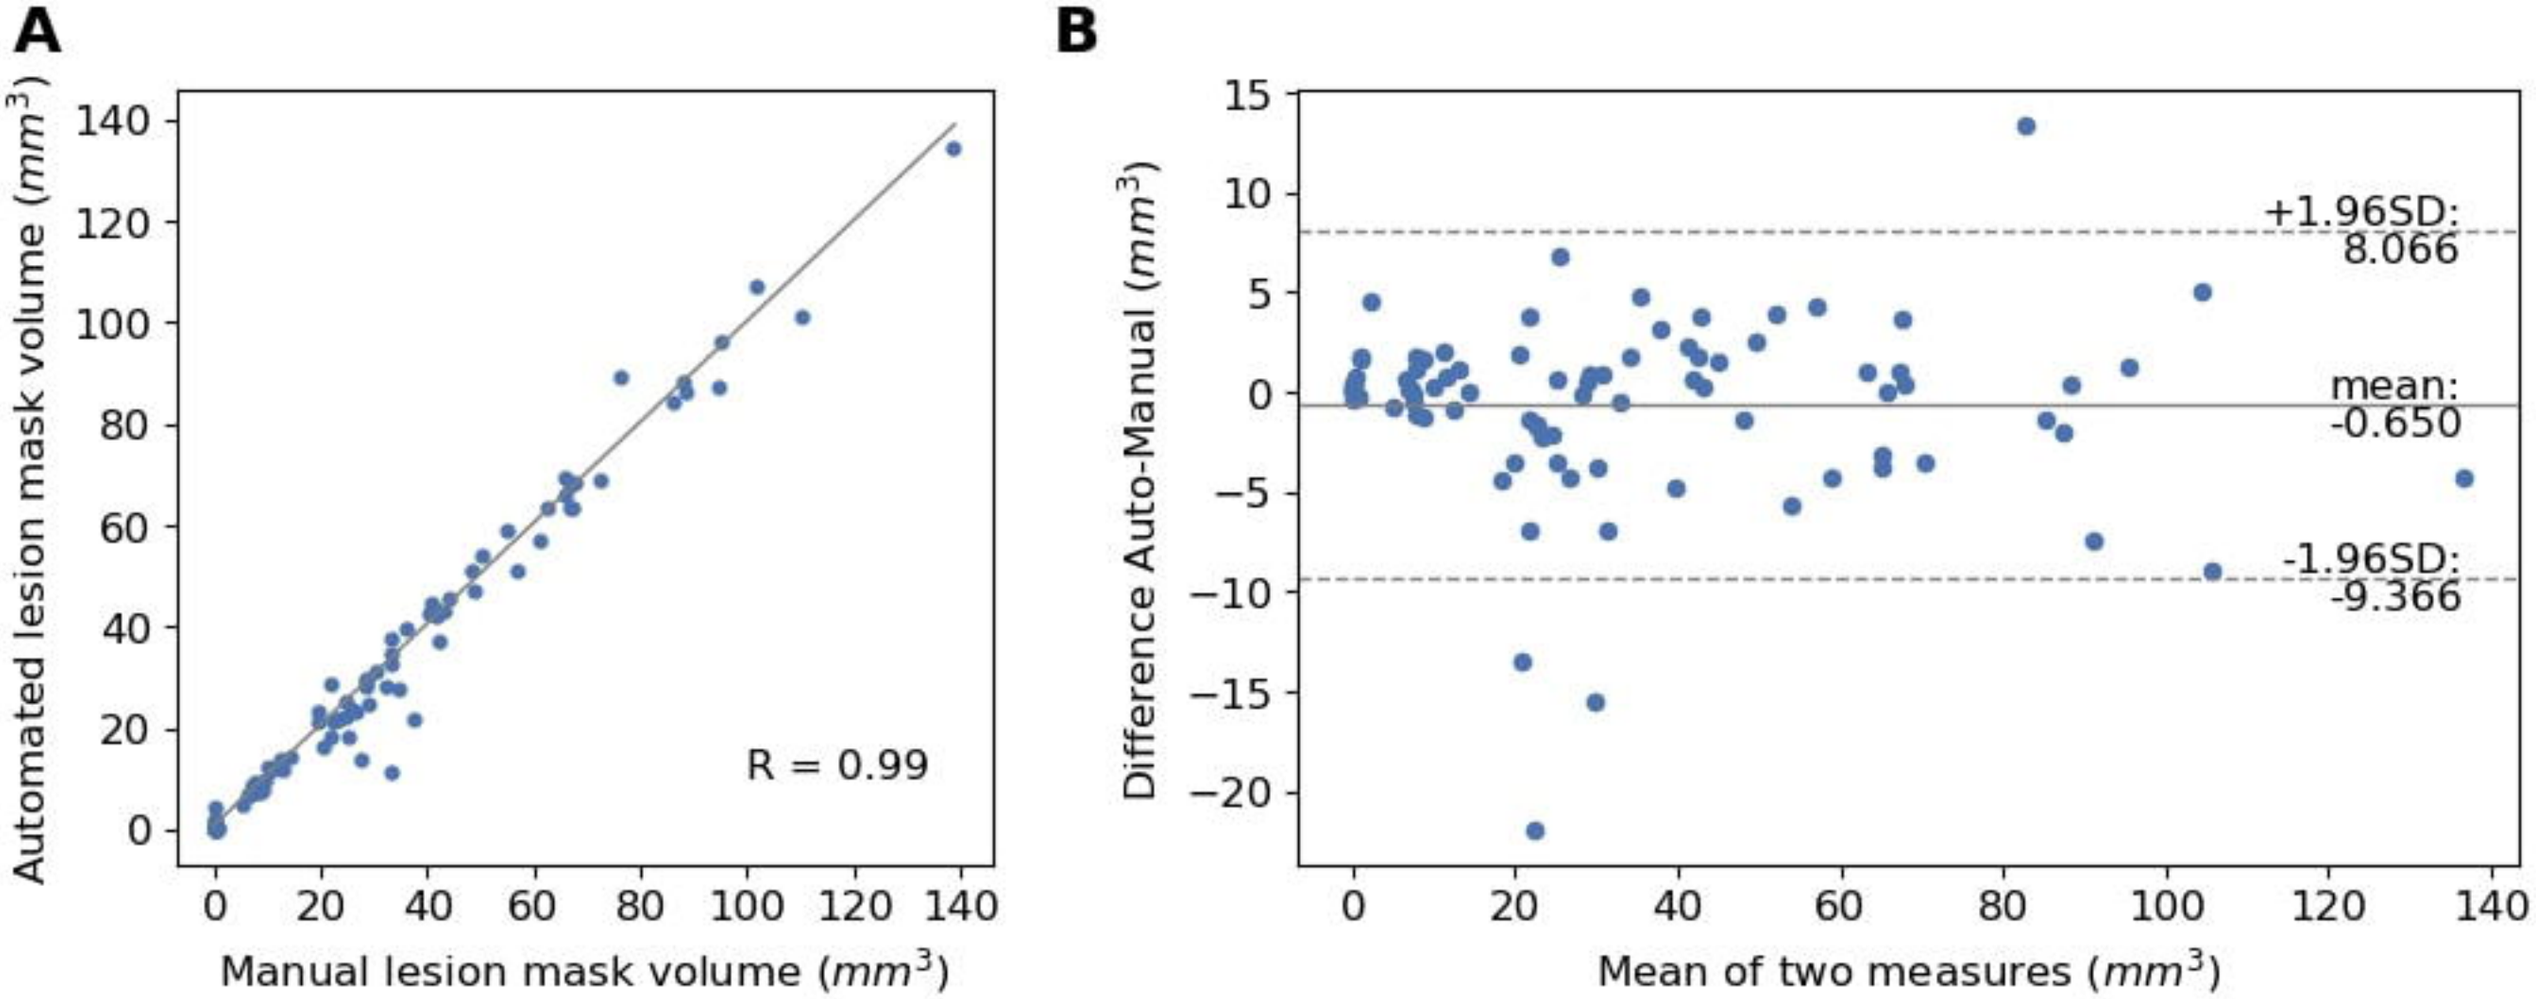
A
Automated lesion mask volume (mm^3)
Manual lesion mask volume (mm^3)
R = 0.99
0
20
40
60
80
100
120
140
B
Difference Auto-Manual (mm^3)
Mean of two measures (mm^3)
+1.96SD:
8.066
mean:
-0.650
-1.96SD:
-9.366
15
10
5
0
−5
−10
−15
−20

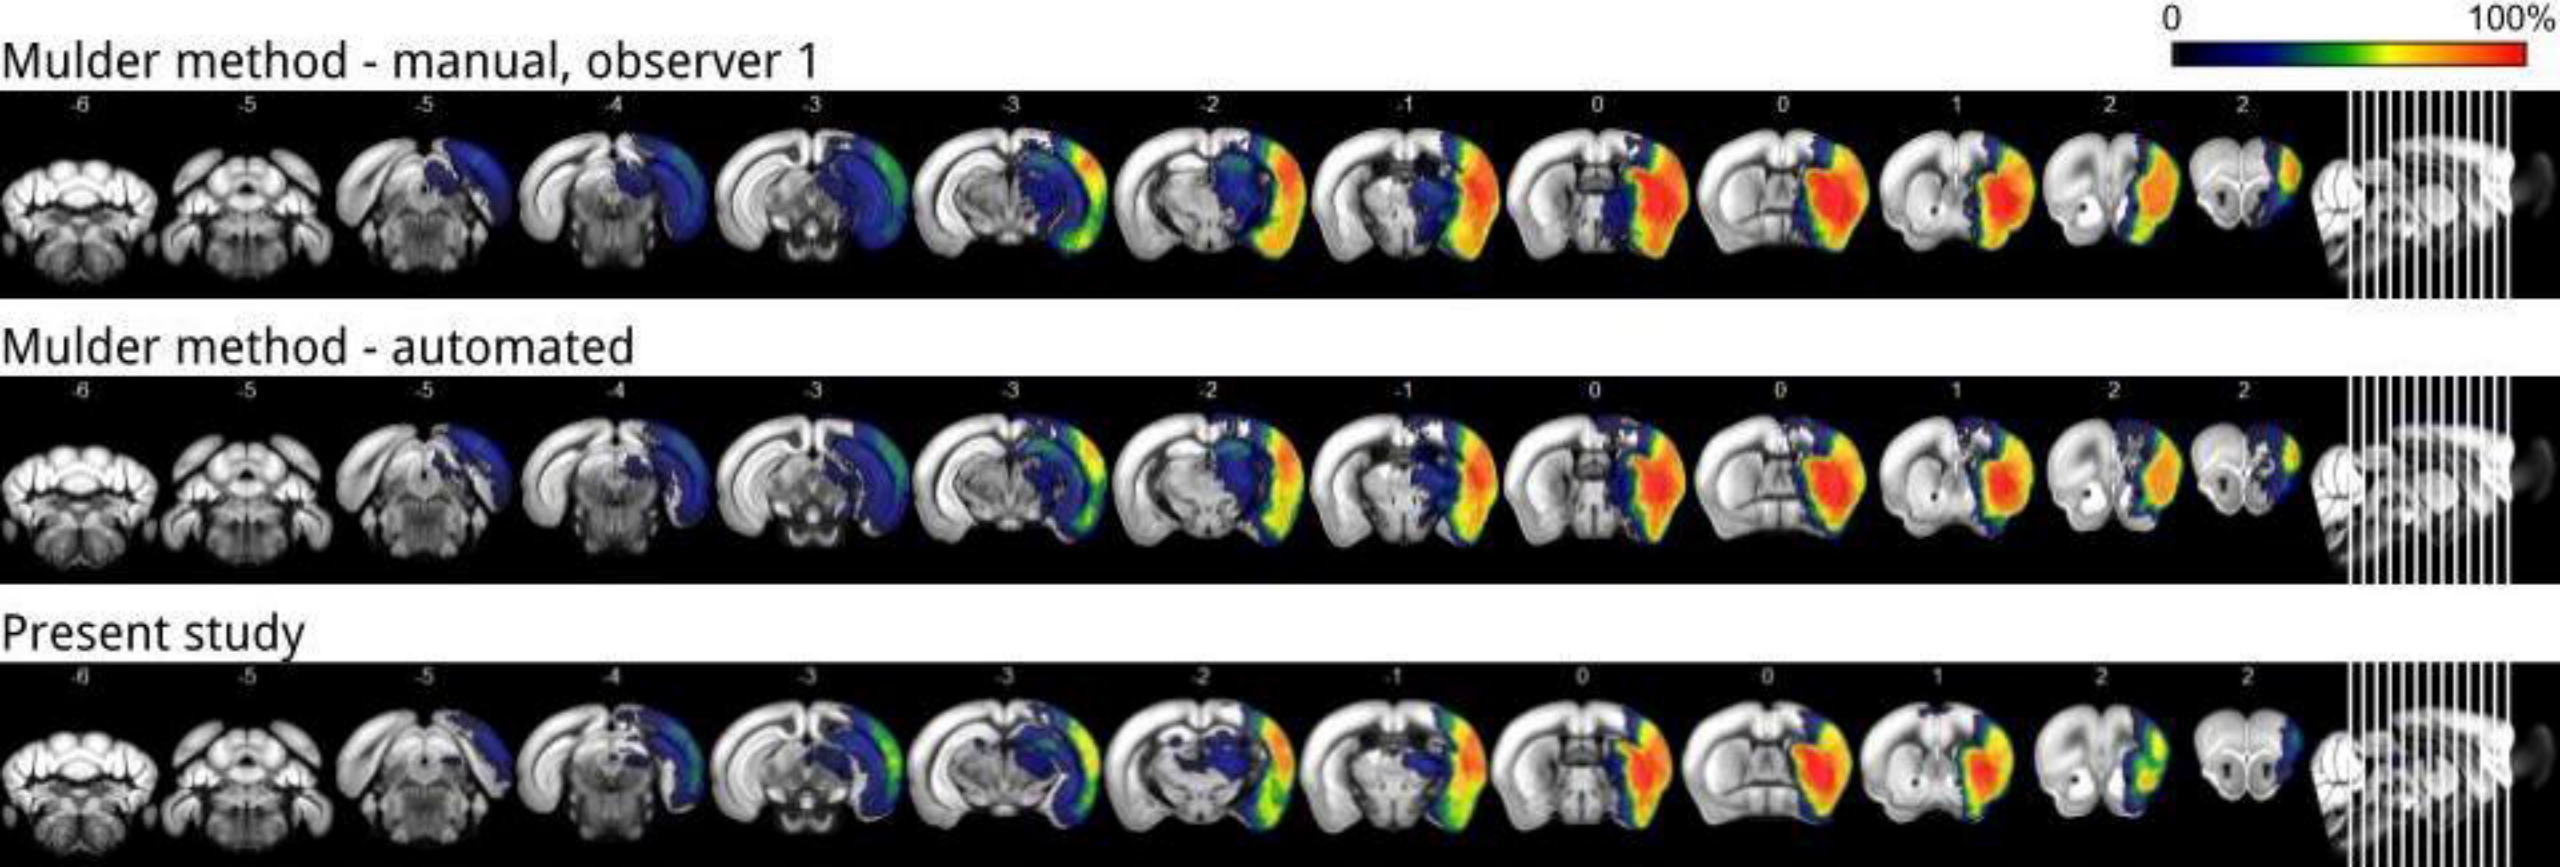
0
100%
Mulder method - manual, observer 1
-6 -5 -5 -4 -3 -3 -2 -1 0 0 1 2 2
Mulder method - automated
-6 -5 -5 -4 -3 -3 -2 -1 0 0 1 2 2
Present study
-6 -5 -5 -4 -3 -3 -2 -1 0 0 1 2 2