



OPEN

## An anomaly detection approach to identify chronic brain infarcts on MRI

Kees M. van Hespen<sup>1</sup>✉, Jaco J. M. Zwanenburg<sup>2</sup>, Jan W. Dankbaar<sup>2</sup>, Mirjam I. Geerlings<sup>3</sup>, Jeroen Hendrikse<sup>2</sup> & Hugo J. Kuijf<sup>4</sup>

The performance of current machine learning methods to detect heterogeneous pathology is limited by the quantity and quality of pathology in medical images. A possible solution is anomaly detection; an approach that can detect all abnormalities by learning how 'normal' tissue looks like. In this work, we propose an anomaly detection method using a neural network architecture for the detection of chronic brain infarcts on brain MR images. The neural network was trained to learn the visual appearance of normal appearing brains of 697 patients. We evaluated its performance on the detection of chronic brain infarcts in 225 patients, which were previously labeled. Our proposed method detected 374 chronic brain infarcts (68% of the total amount of brain infarcts) which represented 97.5% of the total infarct volume. Additionally, 26 new brain infarcts were identified that were originally missed by the radiologist during radiological reading. Our proposed method also detected white matter hyperintensities, anomalous calcifications, and imaging artefacts. This work shows that anomaly detection is a powerful approach for the detection of multiple brain abnormalities, and can potentially be used to improve the radiological workflow efficiency by guiding radiologists to brain anomalies which otherwise remain unnoticed.

In clinical practice, radiologists acquire and assess magnetic resonance (MR) images of the brain for the diagnosis of various brain pathologies. Unfortunately, the process of reading brain MR images is laborious and observer dependent<sup>1–5</sup>. To reduce observer dependence, and to improve workflow efficiency and diagnostic accuracy, automated (machine learning/'artificial intelligence') methods have been proposed to assist the radiologist<sup>6–13</sup>. A common drawback of these methods is their 'point solution' design, in which they are focused on a specific type of brain pathology. Furthermore, the performance of supervised machine learning based solutions is dependent on the quantity and quality of available examples of pathology. In, for example, cerebral small vessel disease, the development of such solutions is challenging, because the parenchymal damage is heterogeneous in image contrast, morphology, and size<sup>14,15</sup>.

A solution that breaks with this conventional approach is anomaly detection: a machine learning approach that can identify all anomalies solely based on features that describe normal data. Because the features of possible anomalies were not learnt, they stand out from the ordinary, and can subsequently be detected. Anomaly detection methods are particularly useful when there is an interest in the detection of anomalous events, but their manifestation is unknown a priori and their occurrence is limited<sup>16,17</sup>. Examples of applications include credit card fraud detection<sup>18</sup>, IT intrusion detection<sup>19</sup>, monitoring of aerospace engines during flight<sup>20</sup>, heart monitoring<sup>21</sup>, detection of illegal objects in airport luggage<sup>22</sup>, or the detection of faulty semiconductor wafers<sup>23</sup>.

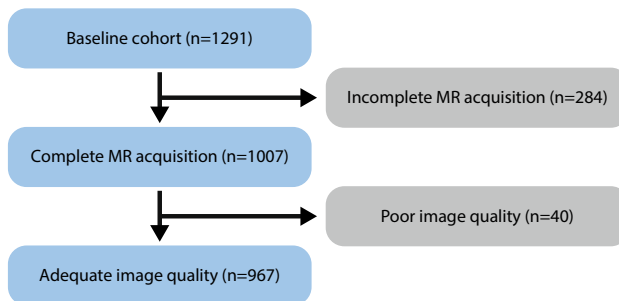
In medical imaging, variational autoencoders and generative adversarial networks have been proposed for anomaly detection tasks. Schlegl et al. have developed a generative adversarial network architecture for the detection of abnormalities on optical coherence tomography images<sup>24,25</sup>. For brain MRI, models have been developed for the detection of tumor tissue<sup>26–28</sup>, white matter hyperintensities<sup>29,30</sup>, multiple sclerosis lesions<sup>31,32</sup>, and acute brain infarcts<sup>28</sup>.

One of the manifestations of cerebral small vessel disease are chronic brain infarcts, including cortical, subcortical, and lacunar infarcts; each with a different appearance on MRI<sup>14,15</sup>. Identification of these infarcts is important, because their occurrence is associated with vascular dementia, Alzheimer's disease, and overall

<sup>1</sup>Center for Image Sciences, University Medical Center Utrecht, Heidelberglaan 100, Postbox 85500, 3584 CX Utrecht, The Netherlands. <sup>2</sup>Department of Radiology, UMC Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands. <sup>3</sup>Julius Center for Health Sciences and Primary Care, UMC Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands. <sup>4</sup>Image Sciences Institute, UMC Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands. ✉email: K.M.vanHespen@umcutrecht.nl

Characteristics	Patients without brain infarcts	Patients with brain infarcts
No. of patients	697	270
Age (years)	57 ( $\pm 10$ )	62 ( $\pm 10$ )
Male sex	550	213

**Table 1.** Demographics of the patient groups with and without brain infarct, used for the training/training-validation, and validation/test sets, respectively.



**Figure 1.** Flowchart showing exclusion of patients in the current study. In total 324 out of 1291 patients were excluded from the current study because of an incomplete MR acquisition (284), where one of two acquired MR images were missing, and because of poor image quality (40).

cognitive decline<sup>33,34</sup>. Because of the heterogeneity in appearance of chronic brain infarcts, location and morphology on MR imaging, anomaly detection would be a possible solution for the identification of these infarcts.

In this study, we constructed an anomaly detection method using a neural network architecture for the detection of chronic brain infarcts from MRI.

## Materials and methods

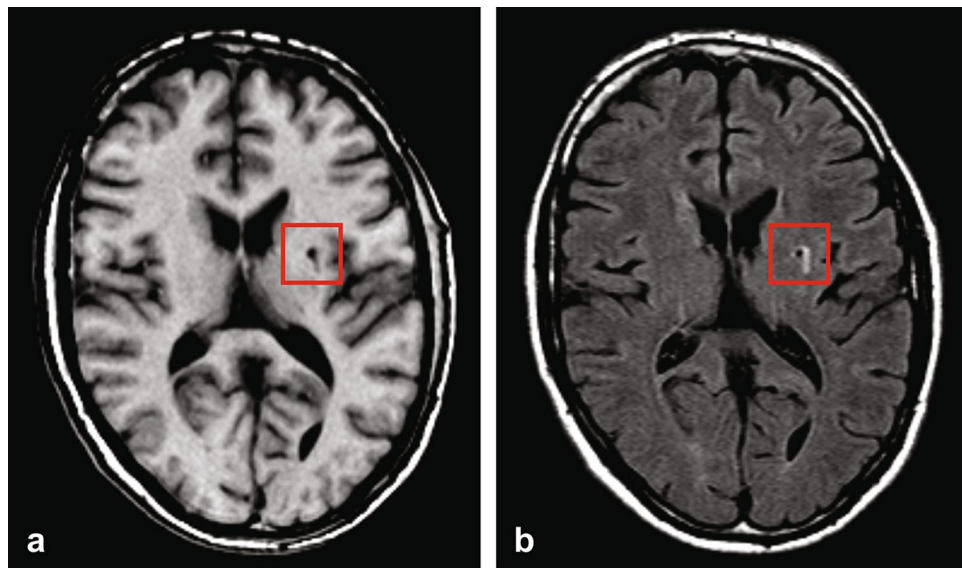
**MR acquisition.** In this retrospective study, we used MR image data from the SMART-MR study<sup>35</sup>, a prospective study on the determinants and course of brain changes on MRI, where all eligible patients that were newly referred to our hospital with manifestations of coronary artery disease, cerebrovascular disease, peripheral arterial disease or an abdominal aortic aneurysm were included after acquiring written informed consent. This study was conducted in accordance with national guidelines and regulations, and has been approved by the University Medical Center Utrecht Medical Ethics Review Committee (METC). In total 967 patients, including 270 patients with brain infarcts (see Table 1 for patient demographics), were included in the current study, see Fig. 1 for exclusion criteria. The imaging data was acquired at 1.5 T (Gyrosan ACS-NT, Philips, Best, the Netherlands), and consisted of a T1-weighted gradient-echo sequence (repetition time (TR) = 235 ms; echo time (TE) = 2 ms), and a T2-weighted fluid-attenuated inversion recovery (T2-FLAIR) sequence (TR = 6000 ms; TE = 100 ms; inversion time: 2000 ms) (example given in Fig. 2). Both MRI sequences had a reconstructed resolution of  $0.9 \times 0.9 \times 4.0 \text{ mm}^3$ , consisted of 38 contiguous transversal slices, and were coregistered<sup>35</sup>. A reasonable request for access to the image data can be sent to Mirjam Geerlings (see author list). The code that supports the findings of this study is publicly available from Bitbucket ([https://bitbucket.org/KeesvanHespen/lesion\\_detection](https://bitbucket.org/KeesvanHespen/lesion_detection)).

All chronic brain infarcts -including cortical infarcts, lacunar infarcts, large subcortical infarct, and infratentorial infarcts- in these images, have been manually delineated by a neuroradiologist with more than 30 years of experience, as described by Geerlings et al.<sup>35</sup> in more detail.

**Image preprocessing.** The images of both acquisitions were preprocessed by applying N4 bias field correction<sup>36</sup>. Additionally, image intensities were normalized such that the 5th percentile of pixel values within an available brain mask was set to zero, and the 95th percentile was set to one.

Two dimensional image patches (smaller subimages of the original image) were sampled within an available brain mask<sup>37</sup> at corresponding locations on both acquisitions for four datasets; a training set, a training-validation set, a validation set, and a test set. For the training set, one million transversal image patches ( $15 \times 15$  voxels) were sampled from images of all but ten patients without brain infarcts. The remaining ten patients without brain infarcts were used for the training-validation set, which was used to assess potential overfitting of the network on the training set. For the training-validation set, 100,000 image patches were randomly sampled. The patches for the training and training-validation set were augmented at each training epoch, by performing random horizontal and vertical flips of the image patches.

The performance of our method on the detection of brain infarcts was evaluated on the validation and test set. The validation set, which was used to evaluate several model design choices consisted of 45 randomly selected



**Figure 2.** Example transversal image slice of the T1-weighted (a) and T2-FLAIR (b) acquisitions. A brain infarct can be observed in the right hemisphere next to the basal ganglia, in the red square, as the hypointense region in the T1-weighted image (a) and the hypointense region with hyperintense ring in the T2-FLAIR image (b).

patients with brain infarcts. The remainder (225) was included in the test set, on which the final network performance was evaluated. In the validation and test sets, 93, and 553 brain infarcts were present, with a median volume of 0.4 ml (range 0.072–282 ml) and 0.44 ml (range 0.036–156 ml), respectively. For these sets, the entire brain was sampled, using a stride of four voxels.

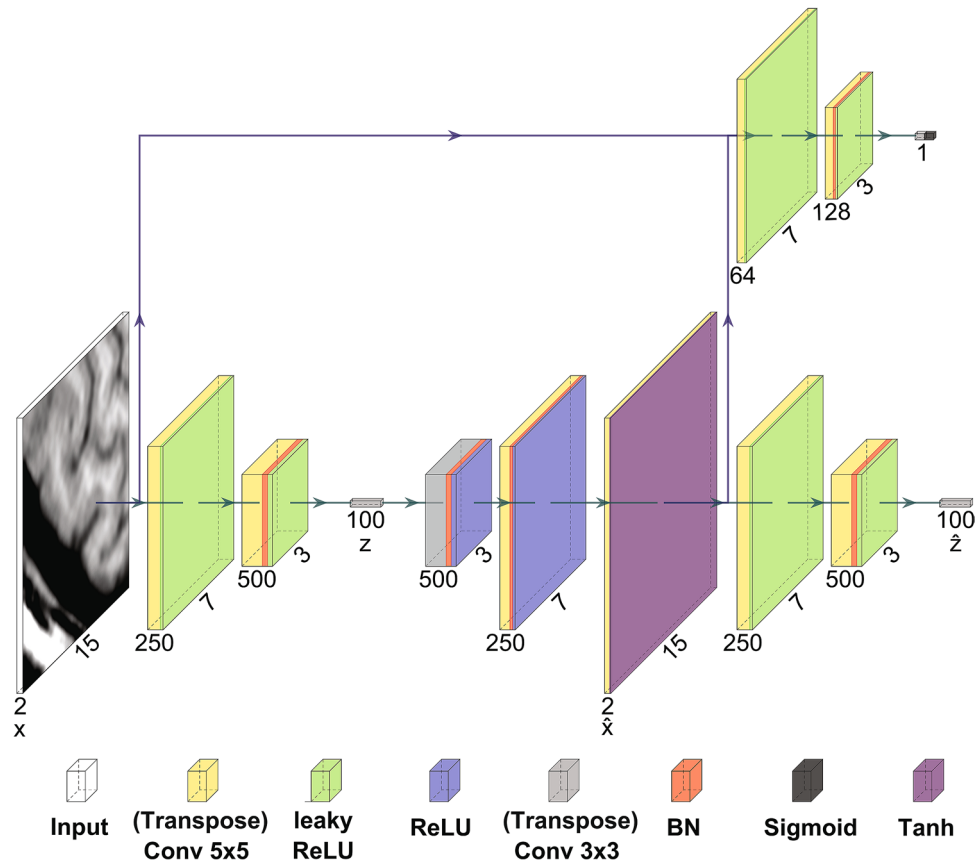
**Network architecture.** We implemented a neural network architecture based on the GANomaly architecture in PyTorch v1.1.0<sup>22,38</sup> which ran on the GPU of a standard workstation (Intel Xeon E-1650v3, 32gb RAM, Nvidia Titan Xp). The neural network (Fig. 3) consisted of a generator (bottom half) and discriminator (top half). The input of the network features two input channels, for both the T1-weighted and the T2-FLAIR image patches. The generator and discriminator consisted of encoder and decoder parts, that each contained three sequential (transposed) convolutional layers, interleaved with (leaky) rectified linear unit (ReLU) activation and batch normalization. The generator was trained to encode the input image patches  $x$  into latent representations:  $z$  and  $\hat{z}$ . Additionally, the generator was trained to realistically reconstruct the input images from the latent vector  $z$  into the reconstructed image  $\hat{x}$ . The discriminator was used to help the generator create realistic reconstructions  $\hat{x}$ . The latent representations  $z$  and  $\hat{z}$  were used to calculate an anomaly score per image patch.

During training of the generator, three error terms were minimized. First, the reconstruction error was computed as the mean difference between the input image and reconstructed image ( $L_{con} = \|x - \hat{x}\|_1$ ). Second, the encoding error was given by the L2 loss between the latent space vectors  $z$  and  $\hat{z}$  ( $L_{enc} = \|z - \hat{z}\|_2$ ). Last, the adversarial loss was computed as the L2 loss between the features from first to last layer of the discriminator, given the input image and reconstructed image ( $L_{adv} = \|f(x) - f(\hat{x})\|_2$ ). To balance the optimization of the network, the generator loss was computed as a weighted combination of the aforementioned losses, with weights of 70, 10, and 1 for  $L_{con}$ ,  $L_{enc}$ , and  $L_{adv}$  respectively.

For the discriminator, two error terms were minimized during training<sup>40</sup>. The first term, is given by the binary cross entropy of the input images and the label of the reconstructed images, and the second term is given by the binary cross entropy of the reconstructed images and the label of the input images. To prevent vanishing gradients in the discriminator, a soft labeling method was chosen, where labels for the reconstructed and input images were uniformly chosen between 0 and 0.2, and between 0.8 and 1, respectively.

The network was trained using the training image patches, which were fed to the network in minibatches of size 64. A learning rate of 0.001 was used, with Adam as optimizer<sup>41</sup>. Training continued until the generator loss on the training-validation set did not decrease any further for ten epochs. The network weights, for the epoch with the lowest generator loss were used for testing.

**Anomaly scoring.** An anomaly score was calculated per image patch as the modified Z-score, a measure of how many median absolute deviations a value lies away from a median value. During training, a median and median absolute deviation were calculated per element of the difference vector  $z - \hat{z}$ , over all training image patches. These values were used to calculate the modified Z-score for all of the difference vector elements of the validation/test image patches. The anomaly score for each image patch was calculated by taking the Nth percentile of modified Z-score values over all vector elements. The value of the percentile was determined in the first experiment.



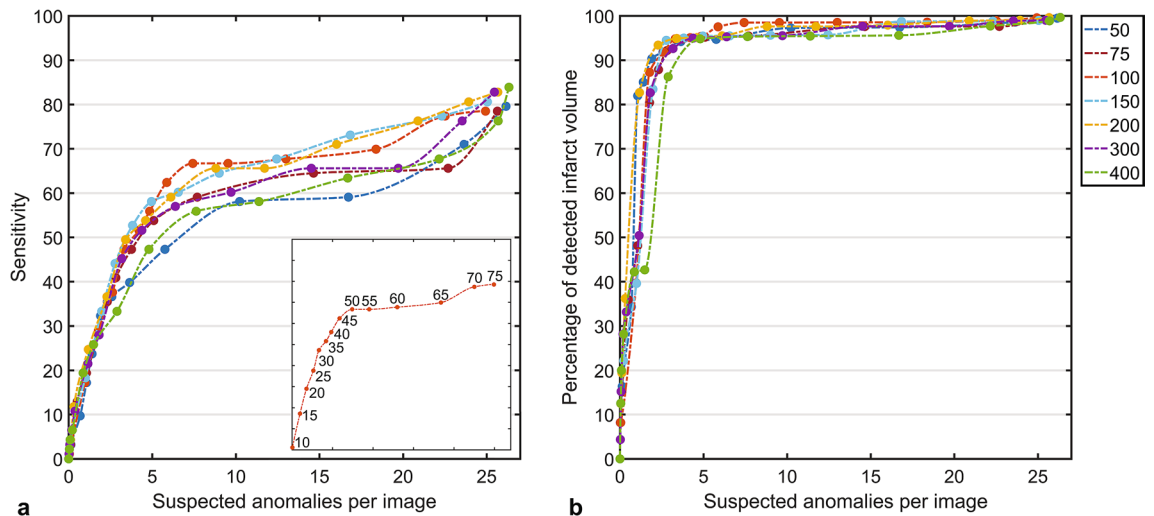
**Figure 3.** Neural network architecture, visualized using adapted software of PlotNeuralNet<sup>39</sup>. The tensor size after each operation is given by the numbers around each box. The input images are encoded into a latent space  $z$ . From the latent space a reconstruction  $\hat{x}$  of the input  $x$  is created. A final encoding of  $\hat{x}$  to  $\hat{z}$  is computed. A discriminator, featured in the top of the image, is fed with images  $x$  and  $\hat{x}$ , where the output of the first to last layer is used for the calculation of the adversarial loss  $L_{adv}$ . Similarly, a reconstruction loss  $L_{con}$  is calculated as the mean difference between  $x$  and  $\hat{x}$ , and an encoding loss  $L_{enc}$  is calculated as the L2 loss between  $z$  and  $\hat{z}$ . The (transposed) conv  $5 \times 5$  layers are applied with a stride of 2 and padding of 1, and are in almost all steps followed by batch normalization (BN). The conv  $3 \times 3$  layer is applied with a stride of 1 and padding equal to 0. The leaky rectified linear unit (ReLU) activation has a negative slope equal to 0.2.

The anomaly scores over all image patches were projected back onto the original brain image, where areas with an anomaly score larger than 3 were flagged as suspected anomalies. Spurious activation, where only a single isolated patch had an anomaly score larger than 3, were filtered out from the final result.

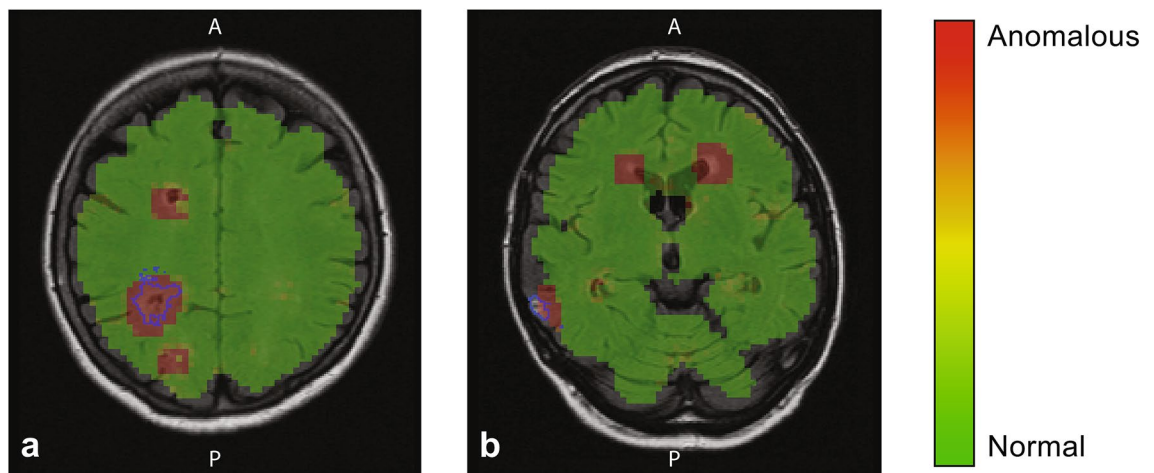
**Experiments.** *Latent vector size and anomaly score calculation.* We investigated the effect of the size of the latent vectors  $z$  and  $\hat{z}$  (Fig. 3), and the effect of the anomaly score calculation on the detection of brain infarcts. We trained several neural network instances with a varying latent vector size, namely: 50, 75, 100, 150, 200, 300, and 400 vector elements. Additionally, we varied the used percentile  $N$  for the anomaly score calculation between 10 and 75, in steps of 5 percentage points. For both parameters, we computed the sensitivity and the average number of suspected anomalies per image, as well as the volume fraction of the detected brain infarcts compared to the total brain infarct volume.

*Suspected anomaly classification.* We used the optimal parameters to evaluate the performance of our proposed method on the test set. We computed the sensitivity and the average number of suspected anomalies per image. In addition, we analyzed the origin of the remaining suspected anomalies, where a neuroradiologist with more than 10 years of experience (JWD), classified these suspected anomalies as one of seven classes, namely: normal tissue, unannotated brain infarct, white matter hyperintensity, blood vessel, calcification, bone and image artefact.

*Missed brain infarcts.* We investigated why our proposed method missed some brain infarcts by evaluating the volume and location of these missed brain infarcts. Additionally, we performed a nearest neighbor analysis, in which we analyzed which training image patches were similar to the test image patches with missed brain infarcts.



**Figure 4.** Brain infarct detection performance for various latent vector sizes and percentile values  $N$ , used in the anomaly score calculation. **(a)** Free Response Operator Curve for all latent vector sizes and **(b)** detected brain infarct volume to total brain infarct volume ratio (given in color). The used percentile ranged between 10 and 75, where for all lines the first datapoint on the left corresponds to 10th percentile, and the right corresponds to the 75th percentile. The inset in **(a)** shows the used percentiles for a latent vector size of 100. The highest sensitivity for the lowest number of suspected anomalies per image (68% and 6) is given by a latent vector size of 100 and a percentile of 50.



**Figure 5.** Anomaly score overlay maps on transversal image slices (T2-FLAIR images). A red color indicates an anomaly score  $\geq 3$ , an anomalous location, where green and yellow an anomaly score  $< 3$ , normal tissue. The manually annotated brain infarcts are given by the blue outlines. An unannotated brain infarct located in a cranial slice can be observed in **(a)**, anterior (A) of the annotated brain infarct. The anomaly at the posterior (P) side of the brain infarct in **(a)** is only delineated in a subsequent image slice. **(b)** Detected cortical brain infarct, and two other suspected anomalous locations, caused by white matter hyperintensities at the horn of the ventricles.

## Results

**Latent vector size and anomaly score calculation.** Based on the validation set, a latent vector size of 100 showed the highest sensitivity and detected brain infarct volume fraction for the same number of suspected anomalies over almost the entire range, compared to all other latent vector sizes (Fig. 4). Similarly, using the 50th percentile yielded an optimal tradeoff between sensitivity and fraction of detected brain infarct volume against the number of suspected anomalies. Given our validation dataset, the optimal parameters for the detection of brain infarcts include the use of the 50th percentile with a latent vector size of 100.

**Suspected anomaly classification.** We used the optimal parameters to evaluate the performance of our proposed method on the test set, on which our proposed method found on average nine suspected anomalies per image (total: 1953, examples are given in Fig. 5). In total, 374 out of 553 brain infarcts were detected by our model (sensitivity: 68%), representing 19.2% of all suspected anomalies (Table 2). These detected brain



Suspected anomaly category	Count (%)
White matter hyperintensity	864 (44.3)
Normal tissue	563 (28.8)
Annotated brain infarcts	374 (19.2)
Image artefact	115 (5.9)
Unannotated brain infarct	26 (1.3)
Blood vessel	6 (0.3)
Calcification	2 (0.1)
Bone	2 (0.1)

**Table 2.** Suspected anomaly categories in the 225 test subjects (Total: 1953. On average: 9 suspected anomalies per image). The suspected anomalies were categorized in eight categories, where the categorization was performed manually by a trained radiologist, except for the ‘annotated brain infarcts’ (these annotations were already available from the used image dataset).

infarcts represented 97.5% of the total brain infarct volume. Eight hundred and sixty-five (44.3%) suspected anomalies were caused by white matter hyperintensities. Image artefacts, e.g. due to patient motion, attributed to 115 (5.9%) of all suspected anomalies. Normal healthy tissue was accountable for 563 (28.8%) of all suspected anomalies. In most cases, these normal tissue false positives were located at tissue—cerebrospinal fluid boundaries. Most interestingly, 26 (1.3%) of all suspected anomalies corresponded to unannotated brain infarcts, which were oftentimes located in the cerebellum and the most cranial image slices.

**Missed brain infarcts.** We performed an additional experiment to better understand why our proposed method missed 179 small brain infarcts (2.5% of the total brain infarct volume). Volume analysis revealed that the missed brain infarcts were in almost all cases smaller than 1 ml (median volume = 0.23 ml). Almost half of the missed brain infarcts (75) were located near the ventricles, at the level of the basal ganglia. Twenty-three missed brain infarcts were present in the cerebellum and 26 in the brain stem. The remaining 55 brain infarcts were located in the cerebral cortex. Automated analysis, where we used a nearest neighbor algorithm to determine which training image patches were similar to the missed brain infarct image patches, revealed that missed brain infarct image patches were oftentimes closely related to training image patches that contained sulci or tissue—cerebrospinal fluid boundaries (see Fig. 6). Furthermore, missed brain infarcts in the brainstem were mostly linked to training image patches in the cortical region, and image patches in the brainstem that border on the cerebrospinal fluid looked similar to cortical image patches with cerebrospinal fluid.

## Discussion

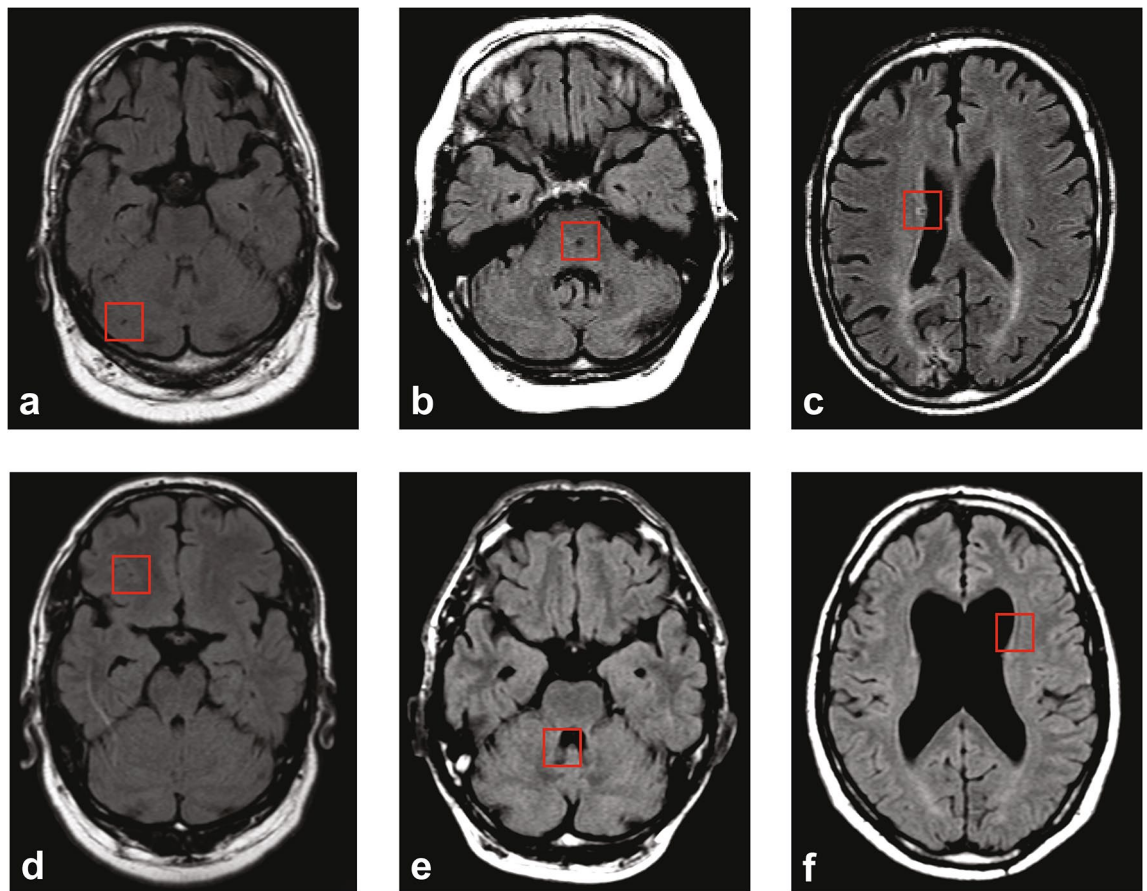
In this paper, we proposed an anomaly detection method for the detection of chronic brain infarcts on brain MR images. Our proposed method detected brain infarcts that accounted for 97.5% of the total brain infarct volume of 225 patients. The missed brain infarcts had in most cases a volume smaller than 1 ml, and were mostly located in the brain stem, and cerebellum, and next to the ventricles. White matter hyperintensities, anomalous calcifications, and imaging artefacts accounted for 44.3%, 0.1%, and 5.9% of suspected anomalies, respectively. Additionally, our proposed method identified additional brain infarcts, which were previously missed by the radiologist during radiological reading (1.3% of all suspected anomalies).

Our proposed method properly evaluated regions that are more easily overlooked by a trained radiologist, given the additional brain infarcts that our method found. The percentage of unannotated brain infarcts that our proposed method found (5% of all brain infarcts) is in accordance with literature, which suggest that reading errors occur in 3–5% of the cases in day-to-day radiological practice. These errors can occur because of an inattentional bias, where radiologists are focused on the center of an image, while overlooking findings at the edges of the acquired image<sup>3</sup>. As workload in radiology is ever increasing due to a larger load of radiological images to be assessed, it is even more easy to overlook brain pathologies<sup>42</sup>. Our presented automated analysis method can potentially alleviate part of this assessment by suggesting anomalous areas for the radiologist to look at, or guiding the radiologist to potentially overlooked brain areas.

Besides suggesting anomalous areas during assessment of the scans by a radiologist, anomaly detection can also be used during image acquisition. In case of an important suspected pathology, this information can then instantly be used to make changes to the acquisition protocol by relocating the field of view or by adding an acquisition that is important for subsequent analysis of the suspected pathology.

Cerebral small vessel disease is a disease with multiple manifestations on MR images. Our proposed method has shown that it can detect at least two of these manifestations, namely chronic brain infarcts -including cortical infarcts, lacunar infarcts, large subcortical infarct, and infratentorial infarcts- and anomalous white matter hyperintensities. This is in contrast to other methods<sup>6,43</sup>, that are only trained for the detection of a single homogeneous type of brain pathology. Similarly, other heterogeneous brain pathologies such as brain tumors, or other manifestations of cerebral small vessel disease can also likely be detected using anomaly detection.

Anomaly detection, which is already used for several years in various fields, such as banking, aerospace, IT, and the manufacturing industry, can also be further explored in the medical imaging field. Potential other applications of anomaly detection include the detection of lung nodules on chest CT images, anomalous regions



**Figure 6.** Comparison between missed brain infarcts and nearest neighbor locations in training images. (a–c) Missed brain infarcts indicated by the red square in the cerebellum, brainstem and cerebrum, respectively, on transversal T2-FLAIR images. (d–f) the corresponding nearest neighbor locations in training data matching (a–c) respectively. The square in (d) contains a part of a normal sulcus that has a similar appearance as the missed brain infarct in the cerebellum in (b).

in the retina, breast cancer in mammograms, calcifications in breast MRI, liver tumor metastasis, or for the detection of areas with low fiber tract integrity in diffusion tensor imaging<sup>25,44–47</sup>. Also for the detection of artefacts in MR spectroscopy or the detection of motion artefacts in MR images, anomaly detection can potentially be beneficial<sup>48,49</sup>. For example, by analyzing the acquired k-space data on motion artefacts during patient scanning, a decision can be made more quickly to redo (parts of) the acquisition.

In future work, the use of 2.5D or 3D contextual information can potentially mitigate problems related to the interpretation of small brain infarcts. This is preferably done with MR images with an isotropic voxel size. This approach would mimic the behavior of human readers who also use contextual information, by scrolling through images, when reading an image. Additional to the performance improvement on small brain infarcts, false positive detections of normal tissue can potentially be mitigated by adding contextual information.

### Limitations

Our method has several limitations. First, its performance on brain infarcts smaller than 1 ml is limited. The detection of small lesions is a common problem that is also present in other medical image analysis applications<sup>50–55</sup>. Ghafoorian et al. have aimed to improve the detection of small lesions by splitting image processing into pathways finetuned to large and small lesions<sup>56</sup>. For our approach, other network architectures might be investigated. Our current architecture seems to predominantly find large anomalies with a relatively large contrast difference compared to the surrounding brain tissue. Other approaches (e.g. a recurrent convolutional neural network<sup>51</sup>) might be able to put more emphasis on finding smaller anomalies with a lower contrast compared to the background. Besides improvements to image analysis, changes in image acquisition may also contribute to better lesion detection. Performance of the detection of small lesions is expected to improve with more up-to-date scanning protocols, which oftentimes have a higher signal-to-noise ratio and/or higher spatial resolution. Van Veluw et al. have shown that microinfarcts can be detected at 7 T MRI<sup>57</sup>, and later work has shown that these small lesions can be imaged at 3 T MRI as well<sup>58</sup>. The current work was done on 1.5 T data with limited sensitivity for detecting small lesions.

Design choices on the neural network architecture were made based on its performance on the validation set that included some brain infarcts. This does not reflect complete anomaly detection, where anomalies should

be completely unknown. However, such an approach is unfeasible in normal practice. Other anomaly detection methods have also used validation sets to tune hyperparameters<sup>59–61</sup>. The use of our validation dataset had no influence on the training of the network, because the network weights of the training epoch with the lowest training-validation loss were used, as opposed to using the network weights of the training epoch with the best brain infarct detection performance on the validation set.

Anomaly detection commonly does not involve classification of anomalies, but only their localization. In case classification is needed, automated methods or manual inspection should be performed after analysis by our proposed method. We envision a workflow in radiology routine where anomaly detection locates possible lesions, a secondary system classifies these lesions or labels them as ‘unknown’, and finally presents the results to a radiologist for inspection.

The method was developed and evaluated on data from a single cohort study. The performance on scans that are acquired on other scanners, from other vendors, and on different field strengths is therefore unknown and a topic of future work. The method could be made applicable to other scanners by partial or full retraining, or by applying transfer learning techniques. In the latter case, a relatively small new dataset might be needed.

The number of false positive detections is relatively high, compared to all other detections (28.8% of all detections). However, we believe that suggesting a few healthy brain locations to the radiologist is less of a problem than missing pathology. Our model has shown its added value by suggesting infarcts that were overlooked by the radiologist.

Lastly, our training set potentially contains unannotated brain infarcts similar to the test set, in which, 5% additional brain infarcts were detected. The potential effect of abnormalities being present in the training data on the final detection performance is likely to be low. The training will be dominated by numerous image patches from normal appearing brain tissue and any abnormalities will therefore have a minimal impact on the anomaly score calculation.

In conclusion, we developed an anomaly detection model for the purpose of detecting chronic brain infarcts on MR images, where our method recovered 97.5% of the total brain infarct volume. Additionally, we showed that our proposed method also finds additional brain abnormalities, some of which were missed by the radiologist. This supports the use of anomaly detection as automated tool for computer aided image analysis.

## Data availability

A reasonable request for access to the image data can be sent to Mirjam Geerlings (see author list). The code that supports the findings of this study is publicly available from Bitbucket ([https://bitbucket.org/KeesvanHespen/lesion\\_detection](https://bitbucket.org/KeesvanHespen/lesion_detection)).

Received: 20 November 2020; Accepted: 23 March 2021

Published online: 08 April 2021

## References

- Hagens, M. H. J. *et al.* Impact of 3 Tesla MRI on interobserver agreement in clinically isolated syndrome: A MAGNIMS multicentre study. *Mult. Scler. J.* **25**, 352–360 (2019).
- Geurts, B. H. J., Andriessen, T. M. J. C., Goraj, B. M. & Vos, P. E. The reliability of magnetic resonance imaging in traumatic brain injury lesion detection. *Brain Inj.* **26**, 1439–1450 (2012).
- Busby, L. P., Courtier, J. L. & Glastonbury, C. M. Bias in radiology: The how and why of misses and misinterpretations. *Radiographics* **38**, 236–247 (2018).
- Brady, A. P. Error and discrepancy in radiology: Inevitable or avoidable?. *Insights Imaging* **8**, 171–182 (2017).
- Lee, C. S., Nagy, P. G., Weaver, S. J. & Newman-Toker, D. E. Cognitive and system factors contributing to diagnostic errors in radiology. *Am. J. Roentgenol.* **201**, 611–617 (2013).
- Guerrero, R. *et al.* White matter hyperintensity and stroke lesion segmentation and differentiation using convolutional neural networks. *NeuroImage. Clin.* **17**, 918–934 (2018).
- Atlason, H. E., Love, A., Sigurdsson, S., Gudnason, V. & Ellingsen, L. M. SegAE: Unsupervised white matter lesion segmentation from brain MRIs using a CNN autoencoder. *NeuroImage Clin.* **24**, 102085 (2019).
- Gabr, R. E. *et al.* Brain and lesion segmentation in multiple sclerosis using fully convolutional neural networks: A large-scale study. *Mult. Scler. J.* **26**, 1217–1226 (2020).
- Devine, J., Sahgal, A., Karam, I. & Martel, A. L. Automated metastatic brain lesion detection: a computer aided diagnostic and clinical research tool. In *Medical Imaging 2016: Computer-Aided Diagnosis* Vol. 9785 (eds Tourassi, G. D. & Armato, S. G.) (International Society for Optics and Photonics, 2016).
- van Wijnen, K. M. H. *et al.* Automated lesion detection by regressing intensity-based distance with a neural network. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019* Vol. 11767 (eds Shen, D. *et al.*) 234–242 (Springer Verlag, 2019).
- Ain, Q., Mehmood, I., Naqi, S. M. & Jaffar, M. A. Bayesian classification using DCT features for brain tumor detection. In *Knowledge-Based and Intelligent Information and Engineering Systems. KES 2010. Lecture Notes in Computer Science* Vol. 6276 (eds Setchi, R. *et al.*) 340–349 (Springer, 2010).
- Shen, S., Szameitat, A. J. & Sterr, A. Detection of infarct lesions from single MRI modality using inconsistency between voxel intensity and spatial location—A 3-D automatic approach. *IEEE Trans. Inf. Technol. Biomed.* **12**, 532–540 (2008).
- Cabezas, M. *et al.* Automatic multiple sclerosis lesion detection in brain MRI by FLAIR thresholding. *Comput. Methods Programs Biomed.* **115**, 147–161 (2014).
- Wardlaw, J. M. *et al.* Neuroimaging standards for research into small vessel disease and its contribution to ageing and neurodegeneration. *Lancet Neurol.* **12**, 822–838 (2013).
- Pantoni, L. Cerebral small vessel disease: From pathogenesis and clinical characteristics to therapeutic challenges. *Lancet Neurol.* **9**, 689–701 (2010).
- Pimentel, M. A. F., Clifton, D. A., Clifton, L. & Tarassenko, L. A review of novelty detection. *Signal Process.* **99**, 215–249 (2014).
- Chandola, V., Banerjee, A. & Kumar, V. Anomaly detection. *ACM Comput. Surv.* **41**, 1–58 (2009).
- Phua, C., Alahakoon, D. & Lee, V. Minority report in fraud detection. *ACM SIGKDD Explor. Newsl.* **6**, 50–59 (2004).
- Jyothsna, V., Prasad, V. V. R. & Prasad, K. M. A review of anomaly based intrusion detection systems. *Int. J. Comput. Appl.* **28**, 26–35 (2011).



20. Clifton, D. A., Bannister, P. R. & Tarassenko, L. A framework for novelty detection in jet engine vibration data. *Key Eng. Mater.* **347**, 305–310 (2007).
21. Lemos, A. P., Tierra-Criollo, C. J. & Caminhas, W. M. ECG anomalies identification using a time series novelty detection technique. In *IFMBE Proceedings* Vol. 18 (eds Müller-Karger, C. et al.) 65–68 (Springer Verlag, 2007).
22. Akcay, S., Atapour-Abarghouei, A. & Breckon, T. P. GANomaly: Semi-supervised anomaly detection via adversarial training. In *Computer Vision - ACCV 2018*, Vol. 11363 LNCS (eds Jawahar, C.V. et al.) 622–637 (Springer International Publishing, 2019).
23. Kim, D., Kang, P., Cho, S., Lee, H. & Doh, S. Machine learning-based novelty detection for faulty wafer detection in semiconductor manufacturing. *Expert Syst. Appl.* **39**, 4075–4083 (2012).
24. Schlegl, T., Seeböck, P., Waldstein, S. M., Schmidt-Erfurth, U. & Langs, G. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *Information Processing in Medical Imaging* (eds Niethammer, M. et al.) 146–157 (Springer International Publishing, 2017).
25. Schlegl, T., Seeböck, P., Waldstein, S. M., Langs, G. & Schmidt-Erfurth, U. f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Med. Image Anal.* **54**, 30–44 (2019).
26. Chen, X. & Konukoglu, E. Unsupervised detection of lesions in brain MRI using constrained adversarial auto-encoders. Preprint at <http://arxiv.org/abs/1806.04972> (2018).
27. Sun, L. et al. An Adversarial Learning Approach to Medical Image Synthesis for Lesion Detection. *IEEE J. Biomed. Heal. Informatics* **24**, 2303–2314 (2020).
28. Alex, V., Safwan, K. P. M., Chennamsetty, S. S. & Krishnamurthi, G. Generative adversarial networks for brain lesion detection. In *Medical Imaging 2017: Image Processing* Vol. 101330G (eds Styner, M. A. & Angelini, E. D.) (International Society for Optics and Photonics, 2017).
29. Bowles, C. et al. Brain lesion segmentation through image synthesis and outlier detection. *NeuroImage Clin.* **16**, 643–658 (2017).
30. Kuijf, H. J. et al. Supervised novelty detection in brain tissue classification with an application to white matter hyperintensities. In *Medical Imaging 2016: Image Processing* Vol. 9784 (eds Styner, M. A. & Angelini, E. D.) (International Society for Optics and Photonics, 2016).
31. Baur, C., Wiestler, B., Albarqouni, S. & Navab, N. Deep autoencoding models for unsupervised anomaly segmentation in Brain MR images. In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* Vol. 11383 LNCS (eds Crimi, A. et al.) 161–169 (Springer Verlag, 2019).
32. Baur, C., Graf, R., Wiestler, B., Albarqouni, S. & Navab, N. SteGANomaly: Inhibiting CycleGAN steganography for unsupervised anomaly detection in brain MRI. In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* Vol. 12262 LNCS (eds Martel, A. L. et al.) 718–727 (Springer Science and Business Media Deutschland GmbH, 2020).
33. van Veluw, S. J. et al. Detection, risk factors, and functional consequences of cerebral microinfarcts. *Lancet Neurol.* **16**, 730–740 (2017).
34. Saczynski, J. S. et al. Cerebral infarcts and cognitive performance. *Stroke* **40**, 677–682 (2009).
35. Geerlings, M. I. et al. Brain volumes and cerebrovascular lesions on MRI in patients with atherosclerotic disease. The SMART-MR study. *Atherosclerosis* **210**, 130–136 (2010).
36. Tustison, N. J. et al. N4ITK: Improved N3 bias correction. *IEEE Trans. Med. Imaging* **29**, 1310–1320 (2010).
37. Anbeek, P., Vincken, K. L., van Bochove, G. S., van Osch, M. J. P. & van der Grond, J. Probabilistic segmentation of brain tissue in MR imaging. *Neuroimage* **27**, 795–804 (2005).
38. Paszke, A. et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* Vol. 32 (eds Wallach, H. et al.) 8024–8035 (Curran Associates, Inc., 2019).
39. Iqbal, H. HarisIqbal88/PlotNeuralNet v1.0.0. (2018). <https://doi.org/10.5281/zenodo.2526396>.
40. Goodfellow, I. J. et al. GAN (generative adversarial nets). *J. Japan Soc. Fuzzy Theory Intell. Inform.* **29**, 177–177 (2017).
41. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. In *Genetic and Evolutionary Computation* (eds Lehman, J. & Stanley, K. O.) 103–110 (ACM Press, 2014).
42. McDonald, R. J. et al. The effects of changes in utilization and technological advancements of cross-sectional imaging on radiologist workload. *Acad. Radiol.* **22**, 1191–1198 (2015).
43. Guo, D. et al. Automated lesion detection on MRI scans using combined unsupervised and supervised methods. *BMC Med. Imaging* **15**, 50 (2015).
44. Zhang, X. et al. Characterization of white matter changes along fibers by automated fiber quantification in the early stages of Alzheimer's disease. *NeuroImage Clin.* **22**, 101723 (2019).
45. Wang, J. et al. Detecting cardiovascular disease from mammograms with deep learning. *IEEE Trans. Med. Imaging* **36**, 1172–1181 (2017).
46. Ouardini, K. et al. Towards practical unsupervised anomaly detection on retinal images. In *Domain Adaptation and Representation Transfer and Medical Image Learning with Less Labels and Imperfect Data* Vol. 11795 (eds Wang, Q. et al.) 225–234 (Springer Verlag, 2019).
47. Tarassenko, L., Hayton, P., Cerneaz, N. & Brady, M. Novelty detection for the identification of masses in mammograms. In *4th International Conference on Artificial Neural Networks*, Vol. 1995, 442–447 (IET, 1995).
48. Kyathanahally, S. P., Döring, A. & Kreis, R. Deep learning approaches for detection and removal of ghosting artifacts in MR spectroscopy. *Magn. Reson. Med.* **80**, 851–863 (2018).
49. Küstner, T. et al. Automated reference-free detection of motion artifacts in magnetic resonance images. *Magn. Reson. Mater. Phys. Biol. Med.* **31**, 243–256 (2018).
50. Kuijf, H. J. et al. Standardized Assessment of automatic segmentation of white matter hyperintensities and results of the WMH segmentation challenge. *IEEE Trans. Med. Imaging* **38**, 2556–2568 (2019).
51. Sudre, C. H. et al. 3D multirater RCNN for multimodal multiclass detection and characterisation of extremely small objects. In *Proceedings of Machine Learning Research* Vol. 102 (eds Cardoso, M. J. et al.) 447–456 (PMLR, 2019).
52. Ngo, D.-K., Tran, M.-T., Kim, S.-H., Yang, H.-J. & Lee, G.-S. Multi-task learning for small brain tumor segmentation from MRI. *Appl. Sci.* **10**, 7790 (2020).
53. Binczyk, F. et al. MiMSeg—An algorithm for automated detection of tumor tissue on NMR apparent diffusion coefficient maps. *Inf. Sci. (Ny)* **384**, 235–248 (2017).
54. Schmidt, P. et al. An automated tool for detection of FLAIR-hyperintense white-matter lesions in Multiple Sclerosis. *Neuroimage* **59**, 3774–3783 (2012).
55. Fartaria, M. J. et al. Automated detection and segmentation of multiple sclerosis lesions using ultra-high-field MP2RAGE. *Invest. Radiol.* **54**, 356–364 (2019).
56. Ghafoorian, M. et al. Automated detection of white matter hyperintensities of all sizes in cerebral small vessel disease. *Med. Phys.* **43**, 6246–6258 (2016).
57. van Veluw, S. J. et al. In vivo detection of cerebral cortical microinfarcts with high-resolution 7T MRI. *J. Cereb. Blood Flow Metab.* **33**, 322–329 (2013).
58. Ferro, D. A., van Veluw, S. J., Koek, H. L., Exalto, L. G. & Biessels, G. J. Cortical cerebral microinfarcts on 3 Tesla MRI in patients with vascular cognitive impairment. *J. Alzheimer's Dis.* **60**, 1443–1450 (2017).

59. Atlason, H. E., Love, A., Sigurdsson, S., Gudnason, V. & Ellingsen, L. M. Unsupervised brain lesion segmentation from MRI using a convolutional autoencoder. Preprint at <http://arxiv.org/abs/1811.09655> (2018).
60. Vasilev, A. *et al.* q-Space novelty detection with variational autoencoders. In *Computational Diffusion MRI*, (eds Bonet-Carne, E. *et al.*) 113–124 (Springer International Publishing, 2020).
61. Alaverdyan, Z., Jung, J., Bouet, R. & Lartizien, C. Regularized siamese neural network for unsupervised outlier detection on brain multiparametric magnetic resonance imaging: Application to epilepsy lesion screening. *Med. Image Anal.* **60**, 101618 (2020).

## Acknowledgements

The Titan Xp used for this research was donated by the NVIDIA Corporation. This work has been made possible by the Dutch Heart Foundation and the Netherlands Organisation for Scientific Research (NWO), as part of their joint strategic research programme: "Earlier recognition of cardiovascular diseases". This project is partially financed by the PPP Allowance made available by Top Sector Life Sciences & Health to the Dutch Heart foundation to stimulate public-private partnerships [Grant numbers 14729 and 2015B028]. J.J.M.Z was funded by the European Research Council under the European Union's Horizon 2020 Programme (H2020)/ERC [Grant number 841865] (SELMA). J.H. was funded by the European Research Council [Grant number 637024] (HEARTOFSTROKE). H.J.K. was funded by ZonMW [Grant number 451001007].

## Author contributions

K.M.H., J.J.Z., J.H., and H.J.K. were involved in the conceptualization of the work. M.I.G. was involved in the acquisition of the study data. Data curation and analysis was performed by K.M.H. and J.W.D. The manuscript was written by K.M.H., J.J.Z., and H.J.K. All authors reviewed and approved the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to K.M.H.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

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



**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/>.

© The Author(s) 2021