MixMicrobleed: Multi-stage detection and segmentation of cerebral microbleeds

Marta Girones Sanguesa[1], Denis Kutnar[1], Bas H.M. van der Velden[1], and Hugo J. Kuijf[1]

[1]Image Sciences Institute, UMC Utrecht

5 August 2021

1 Introduction

Cerebral microbleeds are small, dark, round lesions that can be visualised on T2*-weighted MRI or other sequences sensitive to susceptibility effects [1, 2]. Detection of microbleeds is usually performed visually [3], with the help of validated visual rating scales such as BOMBS [4] or MARS [5]. Semi-automated tools to assist with microbleed detection have been developed in the past [6, 7, 8, 9]. Owing to the blooming effect of microbleeds on MRI, where they appear larger with increasing echo time[1, 10], there have not been many methods focussing on segmentation; since size and volume of microbleeds can change depending on the acquisition settings. Nevertheless, more recent deep learning approaches for microbleed detection, address this as a semantic segmentation task: detection via a method that performs segmentation [11, 12].

In this work, we propose a multi-stage approach to both microbleed detection and segmentation. First, possible microbleed locations are detected with a Mask R-CNN technique [13]. Second, at each possible microbleed location, a simple U-Net [14] performs the final segmentation.

2 Material and methods

2.1 Data

This work used the 72 subjects as training data provided by the Where is VALDO? challenge of MICCAI 2021 (`https://valdo.grand-challenge.org/`). Data consisted of three sequences: T1, T2, and T2*; all aligned in the T2*-space. A binary image including the manual segmentation of microbleeds was provided for every subject.

Data originated from three cohorts and the first number of the subject ID identified the cohort (cohort 1, 2, or 3). The data was split into two separate

datasets, according to the slice thickness of the images. Subjects 1** and 3** had T2* images with 3.0 mm slices; and subjects 2** had images with 0.8 mm slices. Two separate models were trained for these dataset splits.

2.2   Pre-processing

For every patient, image intensities were normalized using the z-score approach. The data originated from different cohorts and all images were resized to a common field-of-view of 512×512 pixels in-plane.

2.3   Mask R-CNN

A pre-trained Mask R-CNN model [13, 15] was finetuned to obtain an initial detection and segmentation of the microbleeds. The method uses 2D patches of size 64×64. Because of the small size of the microbleeds, patches were up-sampled with a factor of four to 256×256. This ensured that the microbleeds had a detectable size in the patches. The three different modalities were introduced as three separate channels. Data augmentation with random affine transformations and horizontal flips was used. The model was trained on 80 % of the data for 15 epochs, with a batch size of 6 and learning rate of 5e-6.

2.4   U-Net

A simple U-Net was applied to obtain the final segmentations. A four-channel input was introduced in this case and consisted of the whole slice in the T2* image, including the previous and posterior slice, and the predicted output of the Mask R-CNN. In case of the first and last slice for every image, previous and posterior slices were blank. Data augmentation was defined by random affine transformations and horizontal flips. The model was trained on 80 % of the data for 50 epochs, with a batch size of 4 and learning rate of 5e-5.

2.5   Post-processing

Because microbleeds consist of dark spots in the T2* images, the first threshold was obtained by determining the minimum intensity value of every microbleed present in the dataset, determined in the T2* images. This value was then maximized and applied to the predictions of both models as a filter.

To further reduce the number of false positives, a mask is applied to the U-Net predicted outputs. Visual inspection showed that most false positives occur at the outer boundaries of the brain. The mask consists of cropping the outside of the brain in the T2* image for every patient and applying a dilation to include the borders of the brain.

Finally, the U-Net output is threshold at 0.001 to obtain a binary mask of the microbleeds present in every image.

2.6 Full prediction pipeline

To summarize, the final prediction in a subject was obtained by following the next steps: First the image intensity is scaled using z-score normalization and both the images and mask are resized to 512×512. A three-channel input is then introduced in the Mask R-CNN to obtain the first prediction, which is thresholded considering the intensity of the T2* image. This prediction is then included in the U-Net input, together with three consecutive slices in the T2* image, as a four-channel tensor. Finally, the U-Net predictions are threshold with the T2* intensities, masked by eliminating the outside and borders of the brain, and threshold to obtain a binary image. Figure 1 shows the full pipeline.
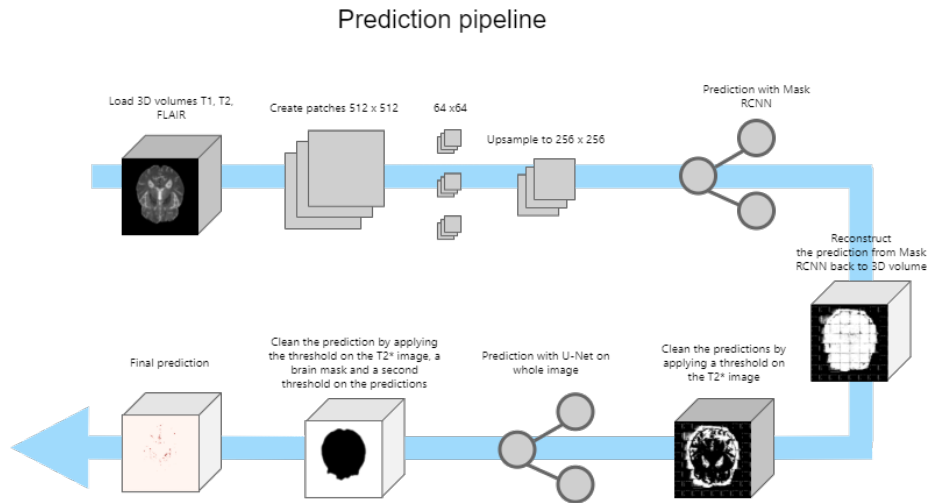

Figure 1: Full microbleed prediction pipeline showing the pre-processing, prediction on the models (Mask R-CNN and U-Net) and post-processing steps to detect and segment the microbleeds in a patient.
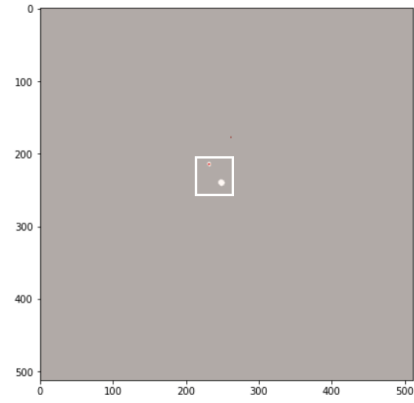
3 Results

Figure 2 shows the confusion matrices of subjects 1**, 2**, and 3**, respectively. Subjects 1** and 3** have been processed with the 'low' slice thickness model, and subjects 2** with the 'high' slice thickness model.
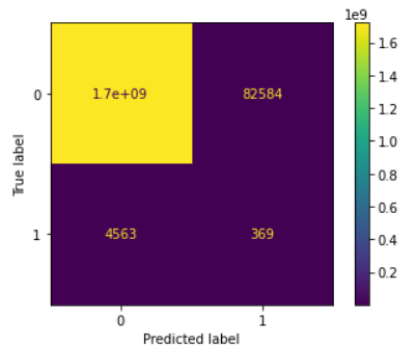
The prediction obtained from a random subject of each of the cohorts has also been included to visualize the output of the segmentation. Note that, to improve the visualization, the white squares show where the true microbleeds are located in the figure.
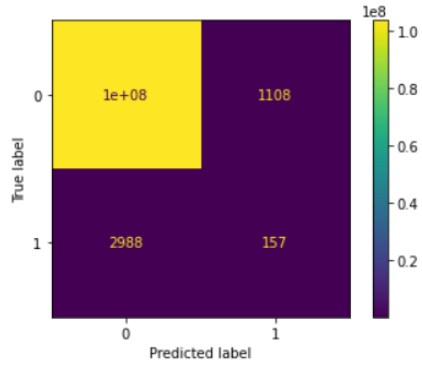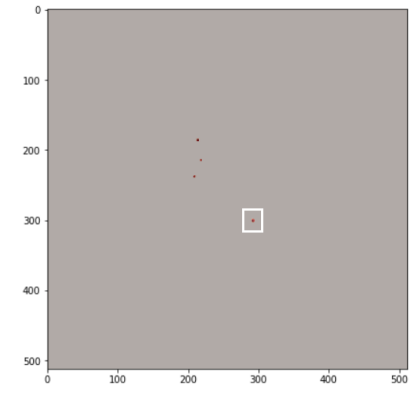
(a) Confusion matrix of subjects 1**.
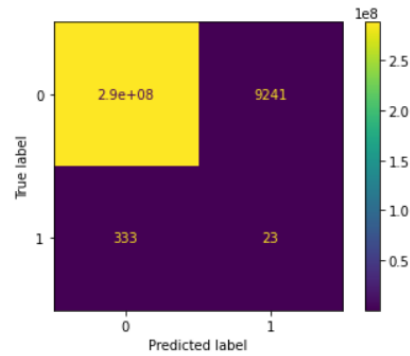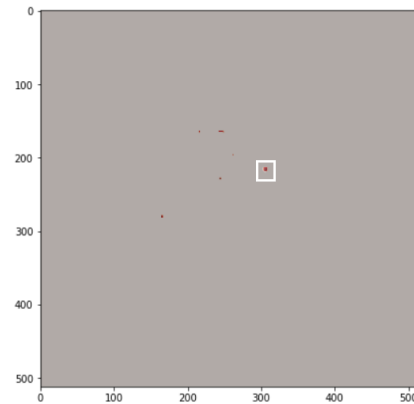

(b) Final segmentation

(c) Confusion matrix of subjects 2**.

(d) Final segmentation

(e) Confusion matrix of subjects 3**.

(f) Final segmentation

Figure 2: Confusion matrices showing the results on the training set of each cohort; and a slice of a random subject showing the true microbleeds (white) and predicted output (red).

4

4   Discussion

Visual inspection of the results on the training data, revealed that most of the false positive detections are dark areas in the image (Figure 3). This corresponds mostly to locations close to the CSF and vessels present in the brain. The use of masks to remove false positive detections was not enough to clear them all, because of the low intensities (similar to true microbleeds) and the central location. A future implementation could use an improved registration and/or segmentation approach to remove the CSF from the images.
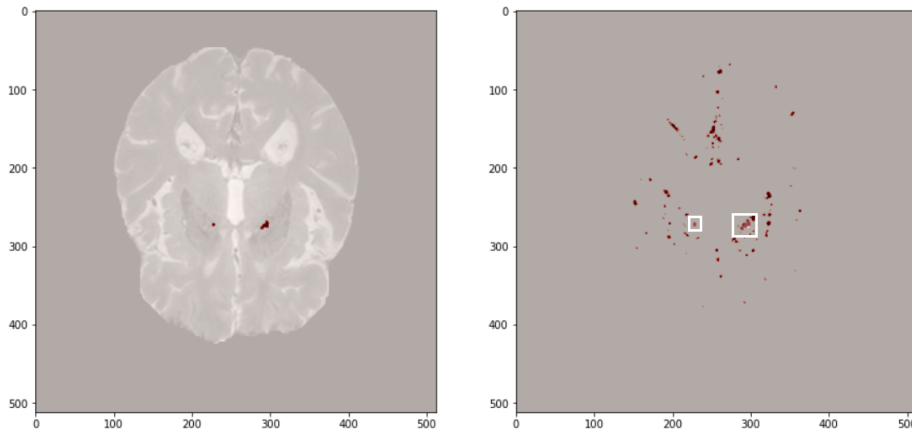

Figure 3: Left: slice of a random subject, showing true microbleeds in red. Right: prediction of the method for this slice (red) with the true microbleeds in the white box. A number of false positives are visible, corresponding to dark regions in the T2* image.

Our current implementation, using a threshold to remove false positive detections, also partially removes the outer boundary of true microbleeds (Figure 3); because the partial volume effect gives it a slightly higher intensity than the signal void at the core of the microbleed. This could be improved by thresholding at the object level (keeping a 3D connected component if at least one voxel survives the intensity threshold) or using a double thresholding and/or region growing approach to retain the borders of true microbleeds.

Similarly, some microbleeds are discard by the last post-processing because they have lower intensities than the threshold, as is shown in the Figure 4.

5   More information

Source code is available at: `https://github.com/hjkuijf/MixMicrobleed`. The docker container hjkuijf/mixmicrobleed can be pulled from `https://hub.docker.com/r/hjkuijf/mixmicrobleed`.
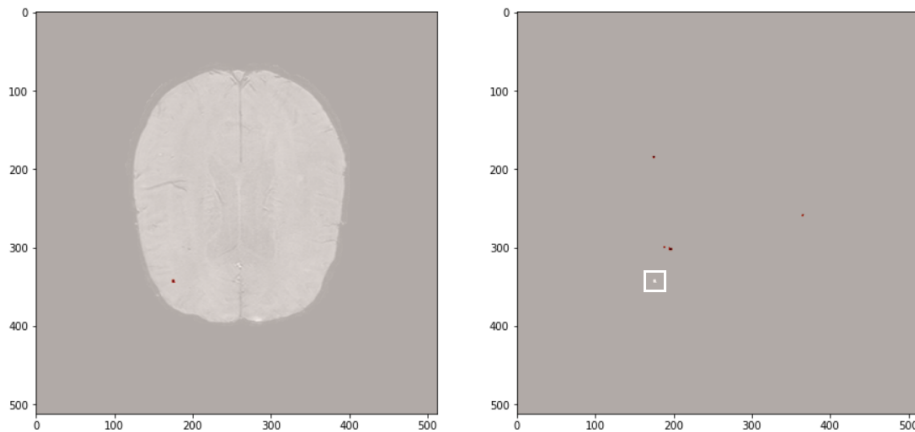
Figure 4: Left: slice of a random subject, showing true microbleeds in red. Right: prediction of the method for this slice (red) with the true microbleeds in the white box. The true microbleed was originally detected, but removed by the post-processing steps.

References

[1] J. M. Wardlaw, E. E. Smith, G. J. Biessels, C. Cordonnier, F. Fazekas, R. Frayne, R. I. Lindley, J. T. O'Brien, F. Barkhof, O. R. Benavente, S. E. Black, C. Brayne, M. Breteler, H. Chabriat, C. Decarli, F.-E. de Leeuw, F. Doubal, M. Duering, N. C. Fox, S. Greenberg, V. Hachinski, I. Kilimann, V. Mok, R. van Oostenbrugge, L. Pantoni, O. Speck, B. C. M. Stephan, S. Teipel, A. Viswanathan, D. Werring, C. Chen, C. Smith, M. van Buchem, B. Norrving, P. B. Gorelick, and M. Dichgans, "Neuroimaging standards for research into small vessel disease and its contribution to ageing and neurodegeneration.," *The Lancet. Neurology*, vol. 12, no. 8, pp. 822–38, 2013.

[2] F. De Guio, E. Jouvent, G. J. Biessels, S. E. Black, C. Brayne, C. Chen, C. Cordonnier, F. E. De Leeuw, M. Dichgans, F. Doubal, M. Duering, C. Dufouil, E. Duzel, F. Fazekas, V. Hachinski, M. A. Ikram, J. Linn, P. M. Matthews, B. Mazoyer, V. Mok, B. Norrving, J. T. O'Brien, L. Pantoni, S. Ropele, P. Sachdev, R. Schmidt, S. Seshadri, E. E. Smith, L. A. Sposato, B. Stephan, R. H. Swartz, C. Tzourio, M. Van Buchem, A. Van Der Lugt, R. Van Oostenbrugge, M. W. Vernooij, A. Viswanathan, D. Werring, F. Wollenweber, J. M. Wardlaw, and H. Chabriat, "Reproducibility and variability of quantitative magnetic resonance imaging markers in cerebral small vessel disease," *Journal of Cerebral Blood Flow and Metabolism*, vol. 36, no. 8, pp. 1319–1337, 2016.

[3] S. M. Greenberg, M. W. Vernooij, C. Cordonnier, A. Viswanathan, R. Al-Shahi Salman, S. Warach, L. J. Launer, M. a. Van Buchem, and M. M.

Breteler, "Cerebral microbleeds: a guide to detection and interpretation.," *Lancet neurology*, vol. 8, pp. 165–74, feb 2009.

[4] C. Cordonnier, G. M. Potter, C. a. Jackson, F. Doubal, S. Keir, C. L. M. Sudlow, J. M. Wardlaw, and R. Al-Shahi Salman, "Improving interrater agreement about brain microbleeds: development of the Brain Observer MicroBleed Scale (BOMBS).," *Stroke; a journal of cerebral circulation*, vol. 40, pp. 94–9, jan 2009.

[5] S. M. Gregoire, U. J. Chaudhary, M. M. Brown, T. a. Yousry, C. Kallis, H. R. Jäger, and D. J. Werring, "The Microbleed Anatomical Rating Scale (MARS): reliability of a tool to map brain microbleeds.," *Neurology*, vol. 73, pp. 1759–66, nov 2009.

[6] M. L. Seghier, M. a. Kolanko, A. P. Leff, H. R. Jäger, S. M. Gregoire, and D. J. Werring, "Microbleed detection using automated segmentation (MIDAS): a new method applicable to standard clinical MR images.," *PloS one*, vol. 6, p. e17547, jan 2011.

[7] S. R. S. Barnes, E. M. Haacke, M. Ayaz, A. S. Boikov, W. Kirsch, and D. Kido, "Semiautomated detection of cerebral microbleeds in magnetic resonance images.," *Magnetic resonance imaging*, vol. 29, pp. 844–52, jul 2011.

[8] H. J. Kuijf, M. Brundel, J. de Bresser, S. J. van Veluw, S. M. Heringa, M. A. Viergever, G. J. Biessels, and K. L. Vincken, "Semi-Automated Detection of Cerebral Microbleeds on 3.0 T MR Images," *PLoS ONE*, vol. 8, no. 6, p. e66610, 2013.

[9] W. Bian, C. P. Hess, S. M. Chang, S. J. Nelson, and J. M. Lupo, "Computer-aided detection of radiation-induced cerebral microbleeds on susceptibility-weighted MR images," *NeuroImage: Clinical*, vol. 2, no. 1, pp. 282–290, 2013.

[10] G. McAuley, M. Schrag, S. Barnes, A. Obenaus, A. Dickson, B. Holshouser, and W. Kirsch, "Iron quantification of microbleeds in postmortem brain.," *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, vol. 65, pp. 1592–601, jun 2011.

[11] Q. Dou, H. Chen, L. Yu, L. Zhao, J. Qin, D. Wang, V. C. Mok, L. Shi, and P.-A. Heng, "Automatic Detection of Cerebral Microbleeds From MR Images via 3D Convolutional Neural Networks," *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1182–1195, 2016.

[12] T. Rashid, A. Abdulkadir, I. M. Nasrallah, J. B. Ware, H. Liu, P. Spincemaille, J. R. Romero, R. N. Bryan, S. R. Heckbert, and M. Habes, "DEEP-MIR: a deep neural network for differential detection of cerebral microbleeds and iron deposits in MRI," *Scientific Reports*, vol. 11, p. 14124, dec 2021.

[13] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, pp. 386–397, mar 2017.

[14] O. Ronneberger, Philipp Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2015. Lecture Notes in Computer Science*, vol. 9351, pp. 234–241, Springer, Cham, 2015.

[15] PyTorch, "TorchVision Object Detection Finetuning Tutorial PyTorch Tutorials 1.9.0+cu102 documentation," 2021.