Deep learning for vessel occlusion classification using CT perfusion maps in acute ischemic stroke

Odysseas Papakyriakou [1,2] O.PAPAKYRIAKOU@AMSTERDAMUMC.NL
Bart J. Emmer [2] B.J.EMMER@AMSTERDAMUMC.NL
Charles B. L. M. Majoie [2] C.B.MAJOIE@AMSTERDAMUMC.NL
Jonathan Coutinho [3] J.COUTINHO@AMSTERDAMUMC.NL
Nerea Arrarte Terreros [1,2,4] N.ARRARTETERREROS@AMSTERDAMUMC.NL
Henk A. Marquering [1,2] H.A.MARQUERING@AMSTERDAMUMC.NL

[1] *Department of Biomedical Engineering and Physics, Amsterdam UMC, The Netherlands*
[2] *Department of Radiology and Nuclear Medicine, Amsterdam UMC, The Netherlands*
[3] *Department of Neurology, Amsterdam UMC, The Netherlands*
[4] *Department of Digital Health & Biomedical Technologies, Vicomtech, Spain*

Abstract

The assessment of the occlusion location in acute ischemic stroke is an essential step in treatment decisions. Occlusions in smaller vessels, like the M2, can be difficult to detect in routine clinical care. Computed tomography perfusion (CTP) maps improve clinicians' accuracy and speed in locating occlusions compared to CT angiography alone. Deep learning (DL) could help automate this process. We propose an attention-based convolutional neural network to classify ICA-T, M1, and M2 occlusions using CTP maps. Our method shows an average accuracy of 79.2% and an F1 score of 86% for M2 occlusions, demonstrating the potential of DL utilizing CTP maps for occlusion location classification.

Keywords: CT perfusion, deep learning, stroke, occlusion location classification

1. Introduction

Acute ischemic stroke (AIS) is one of the leading causes of death and acquired disability, caused by the occlusion of an intracranial artery that interrupts cerebral blood flow. Large vessel occlusions refer to occlusions in proximal intracranial arteries of the anterior circulation, like the intracranial part of the internal carotid artery (ICA) or the M1 segment of the middle cerebral artery (MCA). Medium vessel occlusions (MEVOs) involve smaller and more distal vessels, like the distal M2 and more distal segments of the MCA.

Rapid occlusion detection is necessary to guide AIS treatment, with computed tomography angiography (CTA) and non-contrast computed tomography (NCCT) being the standard imaging modalities. When only CTA is used by clinicians to detect occlusions, MEVOs can be missed by up to 38%, with the majority of these missed cases being M2 occlusions (Duvekot et al., 2021). CT perfusion (CTP) is a dynamic imaging method used to assess various parameters of cerebral perfusion, which are summarized in CTP maps (Christensen and Lansberg, 2019). Studies show that the availability of CTP maps enhances clinicians' accuracy and speed in occlusion detection (Robbe et al., 2024). Beyond improving clinicians' accuracy, CTP maps could also aid automatic vessel occlusion detection. One study used 4D-CTP data for vessel occlusion identification (Bathla et al., 2022), but the use of

CTP maps with DL has not been explored. Yet, CTP maps have been utilized for automatic occlusion classification with an atlas-based method (Peerlings et al., 2023). DL has been shown to outperform atlas-based approaches in medical image analysis tasks (Costea et al., 2022). This work: 1) investigates the use of DL for occlusion classification using CTP maps and 2) develops an attention-based 3D convolutional neural network (CNN) that exploits features from various CTP maps.

2. Methods

Dataset: We include 701 patients from the CLEOPATRA study, which combined AIS patients from one of the CONTRAST multicenter, randomized clinical trials (MR CLEAN-(NO IV, MED, LATE, Registry)) and a local cohort (Koopman et al., 2022). The occlusion location was identified by the CONTRAST imaging core lab and we include patients with one occlusion in: terminus part of ICA (ICA-T), M1, or M2. Our dataset is imbalanced: 14.8% ICA-T, 58.8% M1, 26.4% M2. These proportions are maintained across training, validation, and testing. We use three CTP maps as input channels to the CNN: cerebral blood flow (CBF - millimeters (mL) of blood passing through 100g of brain tissue per minute), cerebral blood volume (CBV - mL of blood in 100g of brain tissue), and time-to-maximum (Tmax - time in seconds at which the contrast tracer reaches its maximum). CBF and CBV reflect relative values in percentages compared to the contralateral hemisphere.

Pre-processing: To bring the CTP maps to a common space, we perform affine registration to an in-house atlas with a centered and straight head using SimpleElastix (Klein et al., 2010). After registration, all maps have an image size of 27 axial planes with 512x512 voxels. We remove artifacts outside the brain using a binary mask computed from the largest connected component, and we clip the values between [0, 20] s for Tmax and between [0, 400] % for relative CBF and CBV. The clipped values are normalized between [0, 1]. A rule-based method determines which hemisphere is affected using the values from the Tmax map: it sums all voxel values on each side of the midline, normalizes the sum by the number of nonzero voxels, and selects the side with the highest value, i.e. longest delay, as the occlusion side. The maps are mirrored across the midline if the occlusion is on the left side, allowing the model to focus solely on learning the discrimination of occlusion locations, without needing to differentiate between the hemispheres.

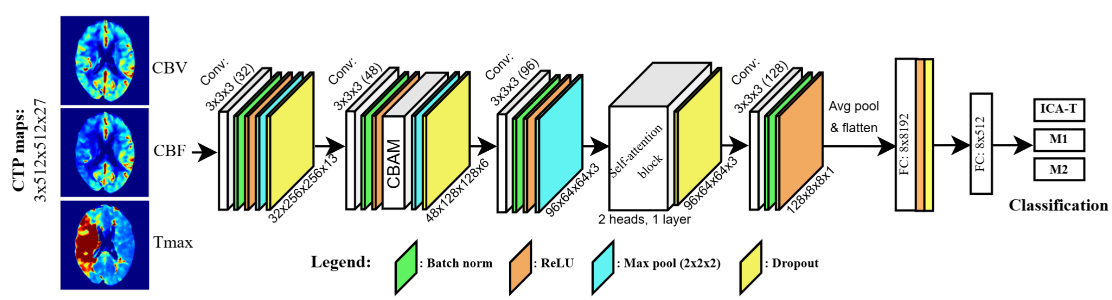

Figure 1: Illustration of our proposed 3D CNN architecture.

Model: We build a 3D CNN that makes use of the convolutional block attention module (CBAM) (Woo et al., 2018). Its channel-wise and spatial attention features allow the model to focus on the most relevant features across the CTP maps and ischemic locations, respectively. To better capture long-range spatial features, the model also includes a self-attention block, as shown in figure 1. To address class imbalance, we use weighted cross-entropy loss with log-smoothed inverse frequency weights and resampling. For training and validation, we use 5-fold cross-validation for 45 epochs on 85% of the dataset with the Adam optimizer, a learning rate of 0.001, and a cosine annealing scheduler. The model with the highest F1 score on the validation set from each fold is used in an ensemble. The ensemble's performance is assessed via voting on the remaining 15% of the dataset for testing.

Evaluation: To evaluate our method, our model is compared against ResNet-CBAM. This architecture combines the ResNet-18 model with CBAM, which has shown promising results in classifying CT scans for stroke or no-stroke (Tahyudin et al., 2024). The ResNet architecture has been used extensively in medical image classification tasks (Xu et al., 2023), as its residual connections allow for meaningful feature extraction (He et al., 2016). CBAM combines channel-wise and spatial attention to capture global contextual information (Woo et al., 2018). The authors integrate these two architectures by applying CBAM on the feature maps generated by ResNet-18, and then use a fully connected network for classification (Tahyudin et al., 2024). Since their method uses two-dimensional CT scans, we implement their architecture as described in their paper, but in three-dimensional space. We train using 5-fold cross-validation as described above, and evaluate the best model of each fold in an ensemble on the same 15% of the dataset on which we evaluate our proposed model.

3. Results

Table 1 presents the confusion matrices for ResNet-CBAM and our proposed occlusion location classification model, using CTP maps for the 106 patients in the test dataset. Both models classify most cases correctly, with ICA-T and M2 occlusions showing the lowest and highest performance, respectively. Our method outperforms ResNet-CBAM for all vessel occlusion locations. For M2 occlusions, 1/28 cases is misclassified as ICA-T, and 3/28 as M1. The overall accuracy of ResNet-CBAM across all occlusion locations is 68.9%, while ours is 79.2%.

		Model output					
		Resnet-CBAM			**Ours**		
		ICA-T	M1	M2	ICA-T	M1	M2
True	ICA-T	9	7	0	11	5	0
	M1	12	44	6	9	49	4
	M2	1	7	20	1	3	24

Table 1: Confusion matrices for Resnet-CBAM and our model for vessel occlusion classification using CTP maps.

Table 2 shows the precision, recall, and F1 score for each occlusion location and for each method. Our method outperforms ResNet-CBAM on all metrics and across all occlusions, classifying correctly 69%, 79%, and 86% of ICA-T, M1, and M2 occlusions respectively.

This is a considerable improvement compared to the 56% for ICA-T and the 71% for M1, and M2 that ResNet-CBAM classifies correctly.

Region	Metric	Resnet-CBAM	Ours
ICA-T	Precision	0.41	0.52
	Recall	0.56	0.69
	F1	0.47	0.59
M1	Precision	0.76	0.86
	Recall	0.71	0.79
	F1	0.73	0.82
M2	Precision	0.77	0.86
	Recall	0.71	0.86
	F1	0.74	0.86

Table 2: Evaluation metrics for each vessel occlusion for Resnet-CBAM and our model.

4. Discussion

This work shows value of DL utilizing CTP maps for vessel occlusion classification in acute ischemic stroke. A previous atlas-based approach using CTP maps (Peerlings et al., 2023) achieved a precision of 57% and a recall of 47% on the M2, whereas our method achieves a precision and recall of 86%. It also outperforms the ResNet-CBAM architecture proposed by Tahyudin et al. (2024) on all metrics and occlusion locations. Automatically classifying M2 occlusions is promising since many M2 occlusions are missed in acute settings (Duvekot et al., 2021). Our model's underperformance for ICA-T occlusion classification could be partially explained by 1) the scarcity of ICA-T occlusions in our dataset and 2) the high variability in downstream hypoperfusion of ICA-T occlusions due to differences in collateral capacity. Commonly, clinicians use CTA to determine the location of the occlusion and the extent of the thrombus. DL methods on vessel occlusion identification have also shown promising results using CTA (Brugnara et al., 2023). As such, future work could combine CTA images with CTP maps to potentially improve occlusion classification accuracy. This work provides primary evidence for the usefulness of CTP maps for vessel occlusion classification with DL, and could potentially help clinicians by classifying the M2 occlusions, which can be difficult to detect in routine clinical care.

References

Girish Bathla, Dhruba Durjoy, Sarv Priya, Edgar Samaniego, and Colin P. Derdeyn. Image level detection of large vessel occlusion on 4D-CTA perfusion data using deep learning in acute stroke. *Journal of Stroke and Cerebrovascular Diseases*, 31, 11 2022. doi: 10.1016/j.jstrokecerebrovasdis.2022.106757.

Gianluca Brugnara, Michael Baumgartner, Edwin David Scholze, et al. Deep-learning based detection of vessel occlusions on CT-angiography in patients with suspected acute ischemic stroke. *Nature Communications*, 14, 12 2023. doi: 10.1038/s41467-023-40564-8.

Soren Christensen and Maarten G. Lansberg. CT perfusion in acute stroke: Practical guidance for implementation in clinical practice. *Journal of Cerebral Blood Flow and Metabolism*, 39:1664–1668, 9 2019. doi: 10.1177/0271678X18805590.

Madalina Costea, Alexandra Zlate, Morgane Durand, et al. Comparison of atlas-based and deep learning methods for organs at risk delineation on head-and-neck CT images using an automated treatment planning system. *Radiotherapy and Oncology*, 177, 11 2022. doi: 10.1016/j.radonc.2022.10.029.

Martijne H. C. Duvekot, Adriaan C. G. M. van Es, Esmee Venema, et al. Accuracy of CTA evaluations in daily clinical practice for large and medium vessel occlusion detection in suspected stroke patients. *European Stroke Journal*, 6:357–366, 12 2021. doi: 10.1177/23969873211058576.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90.

Stefan Klein, Marius Staring, Keelin Murphy, Max A. Viergever, and Josien P. W. Pluim. Elastix : a toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging*, 29(1):196–205, 2010. doi: 10.1109/TMI.2009.2035616.

Miou S. Koopman, Jan W. Hoving, Henk van Voorst, et al. Cost-effectiveness of CT perfusion for patients with acute ischemic stroke (CLEOPATRA)-Study protocol for a healthcare evaluation study. *European Stroke Journal*, 7:188–197, 6 2022. doi: 10.1177/23969873221092535.

Daan Peerlings, Hugo W. A. M. de Jong, Edwing Bennink, et al. Spatial CT perfusion data helpful in automatically locating vessel occlusions for acute ischemic stroke patients. *Frontiers in Neurology*, 14:1136232, March 29 2023. doi: 10.3389/fneur.2023.1136232.

Magretha M. Q. Robbe, Floor M. E. Pinckaers, Bart A. J. M. Wagemans, et al. Additional value of CTP maps in occlusion detection on CTP angiographic reconstructions for ischemic stroke. *European Journal of Radiology*, 181, 12 2024. doi: 10.1016/j.ejrad.2024.111789.

Imam Tahyudin, Anton Satria Prabuwono, Melia Dianingrum, et al. ResNet-CBAM in medical imaging: A high-accuracy tool for stroke detection from CT scans. In *2024 8th*

International Conference on Information Technology, Information Systems and Electrical Engineering, ICITISEE 2024, pages 551–556. Institute of Electrical and Electronics Engineers Inc., 2024. doi: 10.1109/ICITISEE63424.2024.10730079.

Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. CBAM: Convolutional block attention module. In *Computer Vision – ECCV 2018: 15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part VII*, page 3–19, Berlin, Heidelberg, 2018. Springer-Verlag. doi: 10.1007/978-3-030-01234-2_1.

Wanni Xu, You-Lei Fu, and Dongmei Zhu. ResNet and its application to medical image processing: Research progress and challenges. *Computer Methods and Programs in Biomedicine*, 240:107660, 2023. ISSN 0169-2607. doi: https://doi.org/10.1016/j.cmpb.2023.107660.