Thrombus Localization in Acute Ischemic Stroke Using Baseline Lesion Masks

Mahsa Mojtahedi[1] S.M.MOJTAHEDI@AMSTERDAMUMC.NL
Bart Emmer[1] B.J.EMMER@AMSTERDAMUMC.NL
Charles Majoie[1] C.B.MAJOIE@AMSTERDAMUMC.NL
Henk Marquering[1] H.A.MARQUERING@AMSTERDAMUMC.NL
[1] *Amsterdam University Medical Center, 1105 AZ Amsterdam, The Netherlands*

Editors: Under Review for MIDL 2025

Abstract

Thrombus localization is valuable both in treatment selection for ischemic stroke patients and downstream tasks such as extracting thrombus characteristics. We propose the *z-xy* method, a potentially faster alternative to patch-based thrombus localization. *z-xy* consists of two networks: *find-z* which identifies the axial slice containing the thrombus, and *find-xy* that locates the thrombus in that slice. Additionally, we use lesion masks in two ways: (1) to identify and omit the healthy hemisphere; and (2) as potential input to the networks. We compare *z-xy* to a state-of-the-art patch-based method, nnDetection. Both methods are evaluated on: (1) thrombus localization and (2) a downstream task that depends on localized thrombus points - thrombus segmentation. The *z-xy* method performs best without the input lesion mask, with an inference time of 30 seconds per case, an average Euclidean error of 8.2 mm, and a Dice of 0.74 for thrombus segmentation. nnDetection is two times slower than *z-xy* but has an error of 6.3 mm. However, this difference in Euclidean error leads to only a minor improvement in Dice for nnDetection (0.74 ± 0.04 vs 0.77 ± 0.04). Our findings indicate that *z-xy* is a time-efficient alternative to patch-based methods for thrombus localization, particularly in time-sensitive clinical settings of stroke.

Keywords: Localization, Thrombus, Ischemic Stroke

1. Introduction

Ischemic stroke occurs when an artery in the brain is occluded, usually by a thrombus. Stroke can potentially result in severe neurological deficits or even death. Rapidly removing the occluding thrombus is critical, because each passing minute can result in the death of over two million neurons (Desai et al., 2019). Automatically localizing thrombi is useful for two main reasons: first, it aids in stroke diagnosis and treatment management, and second, it improves downstream tasks such as thrombus segmentation by narrowing the region of interest (Mojtahedi et al., 2022). Thrombus segmentation is valuable because it is necessary for calculating imaging characteristics such as thrombus volume and perviousness, which correlate with stroke outcomes (Dutra et al., 2019; Santos et al., 2021).

In order to localize thrombi, spatial context is needed to trace the cerebral arteries and pinpoint the occlusion. However, full high-resolution CTA scans, which are required to detect thrombi, often exceed GPU memory during training. A common solution to this hardware limitation is to use patch-based methods, such as nnDetection (Baumgartner et al., 2021). However, these methods can be slow due to the large number of overlapping patches that must be processed. Here, we propose a potentially faster alternative that uses problem-specific context and ischemic lesion masks to limit the search area.

© 2025 CC-BY 4.0, M. Mojtahedi, B. Emmer, C. Majoie & H. Marquering.

2. Methods and Materials

Our approach outputs the 3D-coordinates of a point on the thrombus. We first identify and omit the healthy hemisphere using automatically generated ischemic lesion masks. Next, the *find-z* network identifies an axial slice containing the thrombus (z-coordinate) on CTA, and the *find-xy* network estimates the thrombus's x and y coordinates within that slice.

Experiments: We compare the accuracy and speed of our approach with nnDetection (Baumgartner et al., 2021), a state-of-the-art medical object localization method.

- *Input ablation:* We investigate if adding ischemic lesion masks as input improves results by training both nnDetection and the *find-z* and *find-xy* networks with two input configurations: (1) CTA scans + lesion masks and (2) CTA scans alone.
- *Downstream task:* To understand the impact of inaccuracies in thrombus localization, we evaluate the resulting point estimations in a downstream task of thrombus segmentation. We use the thrombus segmentation method proposed by Mojtahedi et al. (2022), which requires thrombus points to limit the segmentation region of interest. We compare segmentations generated using thrombus point estimates of each method to those generated based on ground-truth thrombus points.

Data: Our dataset contains 1595 patients with large vessel occlusions who have baseline NCCT and CTA scans from Koopman et al. (2022). The scans are registered to an atlas and resampled to 512 voxels in the sagittal and coronal axes. Ground truth thrombus location points are generated by StrokeViewer, a commercially available software. Lesion masks are created using an in-house-trained nnU-Net that uses co-registered NCCT and CTA scans to generate baseline infarct core and penumbra masks. We use 1350 cases for training, 50 for validation, and 195 for testing, ensuring similar percentages of ICA, M1, and M2 occlusions in each set.

z-xy method: Both *find-z* and *find-xy* networks are 3D ResNet18 architectures. They analyze one axial slice at a time. To provide them with more context, four slices before and after the target slice are also provided as input, resulting in an input size of [Batch size $\times$ 9 $\times$ 512 $\times$ 512]. The *find-z* network estimates the probability of the input slice containing a thrombus. It is trained with binary labels (1 for thrombus, 0 for none) and weighted cross entropy loss ($w_{positive} = 100$). To further alleviate the class imbalance problem, during training, we soften the labels by marking the two slices above and below the ground truth point as positive. During inference, the network analyzes all slices, and the slice with the highest probability is selected. The *find-xy* network estimates the thrombus's (x, y) location coordinates in the selected slice. It is trained using MSE loss.

nnDetection: nnDetection is trained on one fold with the same train/validation split as *z-xy*, using atlas-registered and resampled scans. Ground truth is prepared similar to Brugnara et al. (2023). During inference, the center of the bounding box with the highest confidence is extracted as the thrombus point. We report nnDetection results both with and without test time augmentations (TTA), with the later having a shorter inference time.

Metrics: Inference time per case is the average time for preprocessing and generating outputs across 10 test cases, using a single NVIDIA Tesla V100 GPU with 32 GB memory. Registration and resampling time is not included, as it is used in all methods. Euclidean error refers to: $\sqrt{(x_{gt} - x_e)^2 + (y_{gt} - y_e)^2 + (z_{gt} - z_e)^2}$ (gt: ground truth and e: estimated point). Dice and Surface Dice (Nikolov et al., 2018) are reported for thrombus segmentation.

3. Results

Results in Table 1 show that with TTA (default settings), the best nnDetection (Exp 4) achieves a lower Euclidean error than the best *z-xy* (Exp 1), but it is also about 11 times slower. Interestingly, disabling TTA in nnDetection not only significantly improves its speed but also reduces its error. Adding lesion masks increases the error in the *z-xy* method but provides a slight improvement for nnDetection. The difference in Euclidean error between the best performing *z-xy* (Exp 1) and nnDetection (Exp5) is 1.9 mm. However, even without TTAs, *z-xy* is still twice as fast as nnDetection. Furthermore, in the down-stream task of thrombus segmentation, the higher accuracy of nnDetection results in only 4 % increase in Dice (0.74 ± 0.04 compared to 0.70 ± 0.04).

Table 1: Results of thrombus localization and thrombus segmentation, ± 95% confidence intervals. TTA: test time augmentation, mask: baseline ischemic lesion masks.

Exp	Method	Inputs	Inference Time (min:s) ↓	Euclidean error (mm) ↓	Dice ↑	Surface Dice ↑
1	z-xy	CTA	0:30	8.2 ± 1.1	0.74 ± 0.04	0.80 ± 0.04
2	z-xy	CTA+mask	0:40	11.0 ± 0.9	0.66 ± 0.05	0.72 ± 0.05
3	nnDetection	CTA	4:00	6.9 ± 1.4	0.74 ± 0.05	0.80 ± 0.05
4	nnDetection	CTA+mask	5:33	6.6 ± 1.3	0.75 ± 0.04	0.82 ± 0.04
5	nnDetection no TTA	CTA+mask	1:09	6.3 ± 1.2	0.77 ± 0.04	0.84 ± 0.04

4. Discussion

The *z-xy* method is considerably faster in localizing thrombi than nnDetection. Although *z-xy* results in a larger Euclidean error than nnDetection, this error has minimal impact on the down-stream task of thrombus segmentation. Depending on the application, the speed advantage of *z-xy* may be preferable to the small accuracy improvement offered by nnDetection, making *z-xy* a compelling alternative in time-sensitive stroke settings.

Using lesion masks to limit the area of interest to the affected hemisphere is one of the reasons that *z-xy* is faster than nnDetection. However, adding these lesion masks as input to the networks negatively impacts *z-xy*'s performance. This may be because ischemic lesions form downstream to the occlusion rather than at its exact location. On a brain-wide scale, lesion masks help identify the axial slice containing the occlusion; but within a single slice, they may lack the detail needed for precise x-y localization. In contrast, nnDetection, which uses patches instead of slices, benefits slightly from getting lesion masks as input.

Both methods estimate one point on a 3D object (thrombus). Therefore, a non-zero Euclidean error does not necessarily mean that the estimated point is not on the thrombus. Brugnara et al. (2023) addressed this limitation by creating spheres around the estimated and ground truth points and reporting Dice overlap between them. However, their approach assumes an arbitrary voxel distance from the ground truth as acceptable, and ignores variations in thrombus size and scan spacing. Alternatively, we evaluate the estimated points based on their effectiveness in the downstream task of thrombus segmentation.

The *z-xy* method can assist physicians as a thrombus localization tool, or be part of an automatic thrombus segmentation algorithm that extracts useful thrombus characteristics.

References

Michael Baumgartner, Paul F. Jäger, Fabian Isensee, and Klaus H. Maier-Hein. nnDetection: A Self-configuring Method for Medical Object Detection. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 12905 LNCS, 2021. doi: 10.1007/978-3-030-87240-3{_}51.

Gianluca Brugnara, Michael Baumgartner, Edwin David Scholze, Katerina Deike-Hofmann, Klaus Kades, Jonas Scherer, Stefan Denner, Hagen Meredig, Aditya Rastogi, Mustafa Ahmed Mahmutoglu, Christian Ulfert, Ulf Neuberger, Silvia Schönenberger, Kai Schlamp, Zeynep Bendella, Thomas Pinetz, Carsten Schmeel, Wolfgang Wick, Peter A. Ringleb, Ralf Floca, Markus Möhlenbruch, Alexander Radbruch, Martin Bendszus, Klaus Maier-Hein, and Philipp Vollmuth. Deep-learning based detection of vessel occlusions on CT-angiography in patients with suspected acute ischemic stroke. *Nature Communications*, 14(1):1–15, 8 2023. ISSN 20411723. doi: 10.1038/s41467-023-40564-8. URL https://www.nature.com/articles/s41467-023-40564-8.

Shashvat M. Desai, Marcelo Rocha, Tudor G. Jovin, and Ashutosh P. Jadhav. High Variability in Neuronal Loss: Time Is Brain, Requantified. *Stroke*, 50(1), 2019. ISSN 15244628. doi: 10.1161/STROKEAHA.118.023499.

Bruna G. Dutra, Manon L. Tolhuisen, Heitor C.B.R. Alves, Kilian M. Treurniet, Manon Kappelhof, Albert J. Yoo, Ivo G.H. Jansen, Diederik W.J. Dippel, Wim H. Van Zwam, Robert J. Van Oostenbrugge, Antônio J. da Rocha, Hester F. Lingsma, Aad van der Lugt, Yvo B.W.E.M. Roos, Henk A. Marquering, and Charles B.L.M. Majoie. Thrombus Imaging Characteristics and Outcomes in Acute Ischemic Stroke Patients Undergoing Endovascular Treatment. *Stroke*, 50(8):2057–2064, 8 2019. ISSN 15244628. doi: 10.1161/STROKEAHA.118.024247. URL https://www.ahajournals.org/doi/10.1161/STROKEAHA.118.024247.

Miou S. Koopman, Jan W. Hoving, Henk van Voorst, Jasper D. Daems, Daan Peerlings, Erik Buskens, Hester F. Lingsma, Henk A. Marquering, Hugo W.A.M. de Jong, Olvert A. Berkhemer, Wim H. van Zwam, Marianne A.A. van Walderveen, Ido van den Wijngaard, Aad van der Lugt, Diederik W.J. Dippel, Albert J. Yoo, Bruce C.V. Campbell, Wolfgang G. Kunz, Charles B.L.M. Majoie, and Bart J. Emmer. Cost-effectiveness of CT perfusion for patients with acute ischemic stroke (CLEOPATRA)-Study protocol for a healthcare evaluation study. *European Stroke Journal*, 7(2):188–197, 6 2022. ISSN 23969881. doi: 10.1177/23969873221092535. URL https://pubmed.ncbi.nlm.nih.gov/35647320/.

Mahsa Mojtahedi, Manon Kappelhof, Elena Ponomareva, Manon Tolhuisen, Ivo Jansen, Agnetha A.E. Bruggeman, Bruna G Dutra, Lonneke Yo, Natalie Lecouffe, Jan W Hoving, Henk van Voorst, Josje Brouwer, Nerea Arrarte Terreros, Praneeta Konduri, Frederick J.A. Meijer, Auke Appelman, Kilian M Treurniet, Jonathan M Coutinho, Yvo Roos, Wim van Zwam, Diederik Dippel, Efstratios Gavves, Bart J Emmer, Charles Majoie,

and Henk Marquering. Fully Automated Thrombus Segmentation on CT Images of Patients with Acute Ischemic Stroke. *Diagnostics*, 12(3):698, 3 2022. ISSN 20754418. doi: 10.3390/diagnostics12030698. URL https://www.mdpi.com/2075-4418/12/3/698.

Stanislav Nikolov, Sam Blackwell, Ruheena Mendes, Jeffrey De Fauw, Clemens Meyer, Cían Hughes, Harry Askham, Bernardino Romera-Paredes, Alan Karthikesalingam, Carlton Chu, Dawn Carnell, Cheng Boon, Derek D'Souza, Syed Ali Moinuddin, Kevin Sullivan, DeepMind Radiographer Consortium, Hugh Montgomery, Geraint Rees, Ricky Sharma, Mustafa Suleyman, Trevor Back, Joseph R. Ledsam, and Olaf Ronneberger. Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy. 2018. URL https://arxiv.org/abs/1809.04430.

Emilie M.M. Santos, Nerea Arrarte Terreros, Manon Kappelhof, Jordi Borst, Anna M.M. Boers, Hester F. Lingsma, Olvert A. Berkhemer, Diederik W.J. Dippel, Charles B. Majoie, Henk A. Marquering, and Wiro J. Niessen. Associations of thrombus perviousness derived from entire thrombus segmentation with functional outcome in patients with acute ischemic stroke. *Journal of Biomechanics*, 128:110700, 11 2021. ISSN 0021-9290. doi: 10.1016/J.JBIOMECH.2021.110700.