Cover Page

1) Title of the paper:

Robust Circle of Willis ROI Extraction in MRA via Adaptive Landmark Detection

2) authors' affiliation and address:

LTeN, UMR-6607, Polytech' Nantes, France

&

INSERM, UMR-1087, l'institut du thorax, Nantes, France.

3) e_mail address:

Florent.Autrusseau@univ-nantes.fr

4) Conference & Publisher information:

https://biomedicalimaging.org/2026/

5) bibtex entry:

```
@InProceedings{ISBI_2026,
  author    = {N. Mansouri and R. Bourcier and F. Autrusseau},
  booktitle = {International Symposium on Biomedical Imaging},
  date      = {2026-04-08},
  title     = {Robust Circle of Willis ROI Extraction in MRA via Adaptive
Landmark Detection},
  address   = {London, UK},
  year      = {2026},
}
```

ROBUST CIRCLE OF WILLIS ROI EXTRACTION IN MRA VIA ADAPTIVE LANDMARK DETECTION

N. Mansouri[†], *R. Bourcier*[†], *F. Autrusseau*[⋆†]

[⋆]Polytech'Nantes, LTEN, U-6607, rue Ch. Pauc, 44306, Nantes, France.
[†] Nantes Université, CNRS, INSERM, l'institut du thorax, F-44000 Nantes, France.

ABSTRACT

Accurate localization of the Circle of Willis (CoW) is essential for automated cerebrovascular analysis. However, existing approaches often rely on fixed metadata and struggle when parts of the vessels are not fully visible. We propose a landmark-based ROI localization method that detects key bifurcations and applies learned spatial margins to automatically crop subject-specific ROIs. To achieve full coverage of the CoW, missing bifurcations are estimated using anatomical symmetry or spatial offsets.The method achieved accurate vessel localization and preserved connectivity on two datasets: TopCoW (clDice 0.88, Betti0 0.17) and Lausanne (clDice 0.82, Betti0 0.15). By comparing our approach with the nn-Detection ROI localization, our method provides a more accurate and stable ROI localization. The approach is generalizable and independent of explicit ROI metadata.

Index Terms— Circle of Willis, ROI localization, landmark detection, MRA, cerebrovascular imaging

1. INTRODUCTION

The circle of Willis (CoW) is a critical cerebral arterial network located at the base of the brain that ensures collateral blood flow between the anterior and posterior cerebral circulation [1]. Accurate segmentation and analysis of this structure play a major role in cerebrovascular disease studies like Intracranial aneurysms (ICAs) [2]. However, automated CoW segmentation remains challenging due to anatomical variability, small vessel size, and limited field-of-view in angiographic acquisitions.

Recent advances in deep learning have enabled notable progress in vascular segmentation, yet region-of-interest (ROI) localization remains a critical bottleneck. Several automated methods have been proposed for extracting the Circle of Willis (CoW) prior to segmentation. For example, the eICAB pipeline [3] uses atlas registration and a fixed spherical ROI, followed by extraction of a 90×90×90 voxel patch around the center of mass to standardize network input. TopCoW Challenge strategies define adaptive bounding boxes based on the spatial distribution of vessel voxels [4]. Anatomy-aware frameworks incorporating graph representations infer cerebrovascular network topology and handle missing arteries using connectivity and structural priors [5]. Beyond cerebrovascular imaging, general ROI localization networks combining detection and segmentation have also improved localization accuracy [6]. Despite these advances, most methods still depend on complete anatomy, accurate prior segmentation, or fixed spatial assumptions, limiting their robustness across variable acquisitions and incomplete vascular configurations.

To overcome these limitations, we propose a landmark-based ROI localization method that adapts dynamically to each subject. The approach first detects the main bifurcations of the CoW using a landmark detection model (Nader et al., 2025). We then compute the average spatial margins between bifurcations and ground-truth ROIs across training subjects, and apply these learned margins to new test images. This allows us to automatically crop an ROI centered on the predicted bifurcations. In cases where one or more bifurcations are missing, their locations are estimated using the positions of existing landmarks—either by leveraging anatomical symmetry (e.g., estimating the right MCA from the left) or by applying predefined spatial offsets between related bifurcations (such as between the basilar–PCA and ACA–ACOM). This design makes the ROI estimation both adaptive and robust to incomplete detections, ensuring full coverage of the Circle of Willis even in challenging or asymmetric cases.

2. MATERIAL AND METHODS

2.1. Dataset

For this study, we used the 125 MRA scans released by the TopCoW challenge [7], which includes patients admitted to the Stroke Center at the University Hospital Zurich (USZ) between 2018 and 2019. Although the topCoW dataset contains both MRA and CTA scans, we only used MRA data. The images were acquired on multiple Siemens scanners (1.5T and 3T) in different hospitals in Switzerland, anonymized, and cropped to include only the brain region.

For evaluation, we used two independent datasets to test the robustness of our method. The first set included 17 MRA cases from TopCoW, held out from training. The second set comprised 20 MRA images from the OpenNeuro Lausanne TOF-MRA Aneurysm Cohort [8]. The ground truth ROIs for both datasets were provided, allowing for quantitative evaluation. These datasets enabled comprehensive testing of our ROI localization approach across different imaging conditions and anatomical variations.

2.2. Landmark Detection

Based on the approach proposed by Nader et al. [9], CoW bifurcations are detected using a two-step landmark prediction framework. In the first step, approximate landmark regions are identified using nn-Detection [10]. In the second step, the precise coordinates of each bifurcation within these regions are predicted using an encoder–decoder network that performs heatmap regression. This two-stage design ensures accurate and robust localization of key bifurcations, which we then use as reference points for our adaptive ROI extraction.

The landmark detection algorithm identifies key bifurcation points (Fig. 1) including the anterior cerebral artery (ACA), anterior communicating artery (ACOM), middle cerebral artery (MCA), internal carotid artery (ICA), posterior cerebral artery (PCA), and basilar artery (BA). We chose these seven major bifurcations among others because they lie close to the edges of the ROI, which make them ideal reference points for defining its boundaries. Using their predicted 3D coordinates, we can adaptively crop the ROI to capture the relevant vascular segments around each bifurcation.

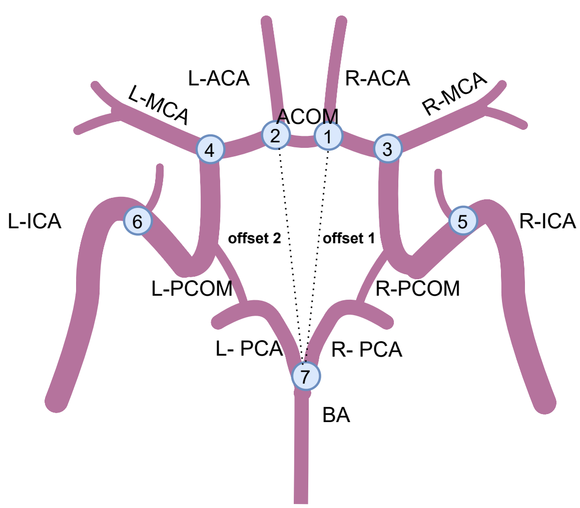

Fig. 1. The Circle of Willis geometry and corresponding bifurcation landmarks as predicted by the automated detection model.

2.3. ROI Estimation & Cropping

To define the optimal cropping region around the Circle of Willis, we first measured the distances between each bifurcation and the edges of its annotated ROI in all six spatial directions ($\pm x$, $\pm y$, $\pm z$) across all training images. These distances capture how far each bifurcation typically lies from the ROI boundaries. We then averaged these distances across all 108 training images to compute a set of learned margins, representing the typical anatomical extent of the Circle of Willis.

During inference, the predicted bifurcations for a new image are used to generate an initial ROI. The learned margins are then added to this ROI in each direction to slightly enlarge the box, ensuring that all arteries are fully captured. The final ROI is therefore the smallest 3D region that encloses all detected bifurcations, extended by the learned margins (Fig. 2-b). This approach allows automatic, subject-specific cropping while preserving the full vascular structure.

2.3.1. Handling missing bifurcations

In instances where certain bifurcations were not detected by the landmark model, we employed estimation strategies based on anatomical symmetry and learned spatial relationships. For lateral bifurcations (e.g., right or left MCA), missing landmarks were mirrored across the mid-sagittal plane using the contralateral coordinates. For bifurcation 13, which connects the BA to the PCAs, we first attempted to estimate its position using learned offsets from other landmarks (1 or 2, corresponding to the ACA–ACOM connections) (Fig. 1). If no learned offset was available, a fallback mirroring along the head–foot (z) axis was applied. These procedures ensured that all key bifurcations required for ROI cropping were available.

2.4. Evaluation Metrics

The proposed ROI prediction approach was evaluated using geometric and topological metrics. The centerline Dice coefficient (clDice) measured the overlap between predicted and reference vessel centerlines, while the Betti0 error quantified differences in the number of connected components. Evaluation was performed by comparing predicted ROIs with manual annotations, and reported values represent the average performance across all test cases in each dataset.

3. ROI PREDICTION RESULTS

The proposed ROI prediction approach demonstrated robust and consistent performance across datasets. On the TopCoW data, the model achieved an average clDice of 0.88, indicating excellent alignment between the predicted and reference centerlines, and an average Betti0 error of 0.17, confirming that the vascular connectivity was well preserved. Similarly, on the Lausanne dataset, the method achieved an average clDice

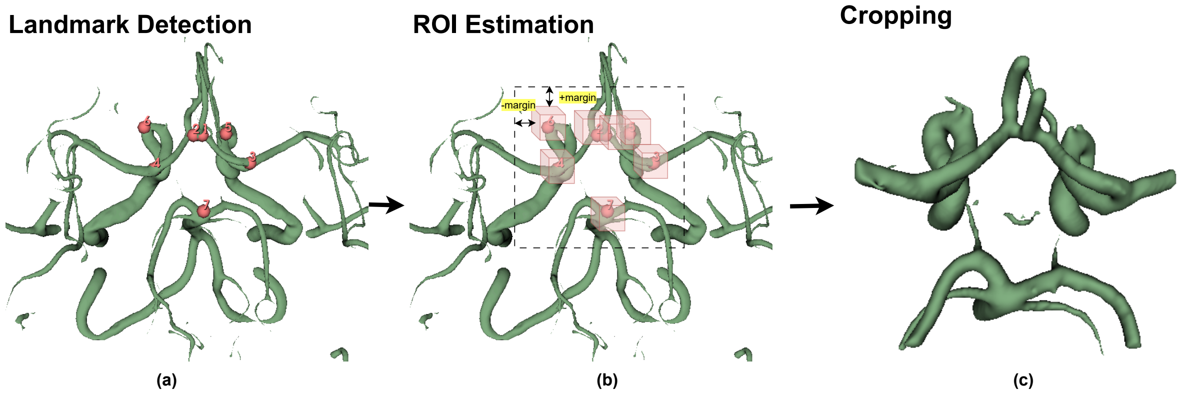

Fig. 2. Overview of the ROI prediction process. (a) Example of an MRA image with bifurcations automatically detected using the landmark detection model. (b) ROI estimation obtained by applying the calculated margin around the detected bifurcations, shown as red bounding boxes. (c) Final cropped ROI, defined as the smallest 3D region enclosing all detected bifurcations.

of 0.82 and an average Betti0 error of 0.15, showing that the model maintained accurate vessel localization and preserved the overall vascular topology despite anatomical variability and differences in image acquisition. These results highlight the robustness and generalizability of the proposed ROI prediction strategy for automatic Circle of Willis localization, even in the absence of explicit metadata on ROI position or size.

Our method was compared with an ROI localization approach based on nn-Detection, which was used by one of the top-performing teams in the TopCoW challenge [7]. To evaluate accuracy, we measured the distances between the predicted center-of-mass (CoM) and the ground truth (GT) for 17 cases. For reference, we also computed the CoM distances using the nn-Detection approach. The results are summarized in Figure 3. As shown in the plot, the distances between the GT CoM and our predicted CoM are consistently smaller than those obtained with nn-Detection. Across all cases, our method remains within a tight range of approximately 0.9 to 4.6 mm (mean $\approx$ 2.4 mm), indicating high spatial precision and stable localization of the ROI.

In contrast, nn-Detection exhibits much larger errors, with distances spanning from about 5 mm up to nearly 38 mm (mean $\approx$ 17.3 mm). This variability reflects a reduced robustness, particularly in challenging cases where the GT position varies or vessels are only partially visible. Overall, our approach provides a markedly more accurate and reliable CoM estimation than nn-Detection. Similar performance was obtained on the Lausanne dataset.

Moreover,our method using landmark detection is considerably faster than the nnDetection apporach. This is expected because nnDetection requires a full 3D convolutional network pass over the entire volume, which is computationally intensive, whereas our approach directly predicts key bifurcation

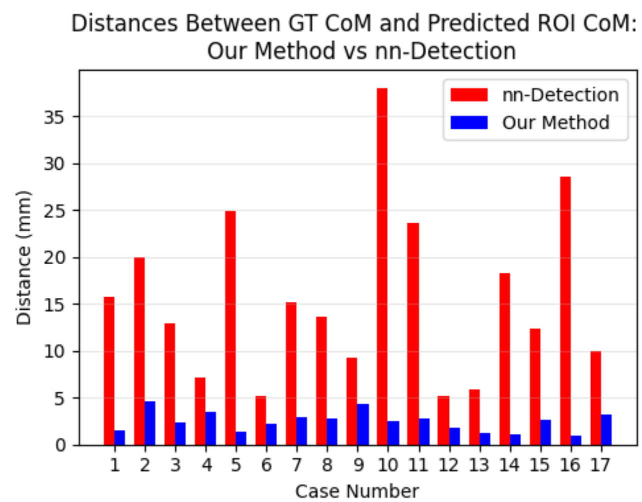

Fig. 3. Comparison of ROI center-of-mass distances. The plot shows the distances between the ground truth (GT) center of mass and the predicted ROI CoM for both our method and the nn-Detection approach on the TopCoW test set.

points, reducing the amount of data processed and the complexity of inference.

3.1. Qualitative Evaluation

We conducted a qualitative evaluation to illustrate the performance of the ROI prediction and bifurcation estimation process (Fig.4). In a TopCoW case, the landmark detection model initially missed bifurcations 1 and 2, located between the ACA and ACoM on the right and left hemispheres, while detecting the other bifurcations (shown in orange). Using the average spatial offsets computed from the training data, the

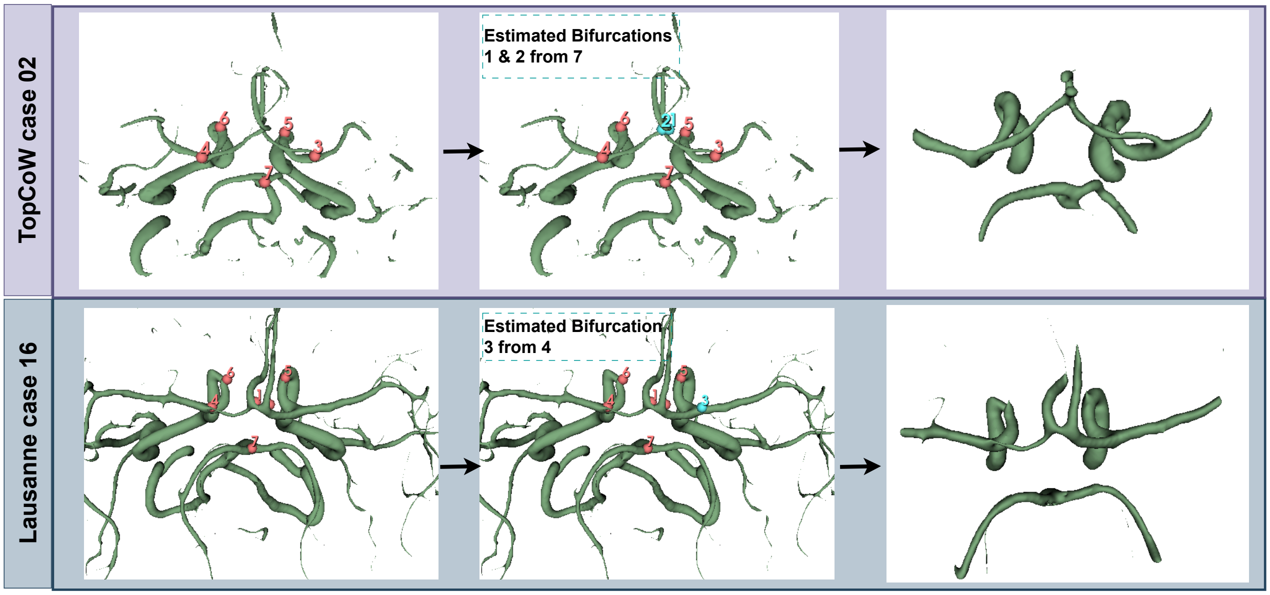

Fig. 4. Qualitative evaluation of ROI prediction and bifurcation estimation. Top row : TopCoW case. First image : all detected bifurcations in orange, with bifurcations 1 and 2 missing. Second image : missing bifurcations 1 and 2 estimated in blue using average offsets from the training set, alongside previously detected ones. Third image : final ROI capturing all CoW arteries. Bottom row : Lausanne case 16. First image : detected bifurcations in orange, with bifurcation 3 missing. Second image : bifurcation 3 estimated in blue using anatomical symmetry from bifurcation 4. Third image : final ROI including all relevant arteries.

model successfully estimated the missing bifurcations (shown in blue), allowing the definition of a final ROI that captured all major CoW arteries. In a Lausanne case (case 16), bifurcation 3, situated between the right ICA and the right MCA, was initially missed. Its location was estimated using anatomical symmetry from bifurcation 4, producing a complete bifurcation prediction (blue) and a final ROI that encompassed the entire CoW. These examples demonstrate that the proposed approach can reliably handle incomplete detections, adaptively estimating missing bifurcations to ensure robust and comprehensive ROI extraction across different datasets.

4. DISCUSSION & CONCLUSION

Our results demonstrate that the proposed landmark-based ROI prediction method is both robust and generalizable across different datasets. By dynamically estimating missing bifurcations using average training offsets or anatomical symmetry, the approach ensures full coverage of the Circle of Willis even in cases with incomplete detections. Quantitative evaluation showed high clDice scores and low Betti0 errors, confirming accurate localization and preservation of vascular connectivity. Qualitative assessment further illustrated that the method reliably captures all major bifurcations and produces precise ROIs.

Furthermore, when compared to the nn-Detection ROI localization approach, our method consistently produces predicted centers of mass that are closer to the ground truth. While nn-Detection shows large variability in some cases, our approach maintains stable and accurate localization across all samples. This comparison highlights the improved precision and robustness of our landmark-based method, making it particularly suitable for automated cerebrovascular analysis where accurate ROI definition is critical.

5. ACKNOWLEDGMENTS

This work was partially supported by the French RHU-ANR project “eCAN” (#ANR-23-RHUS-0013), and PHC Toubkal 23/162.

6. REFERENCES

[1] Zvonimir Vrselja, Hrvoje Brkic, Stefan Mrdenovic, Radivoje Radic, and Goran Curic, “Function of circle of willis,” *Journal of Cerebral Blood Flow & Metabolism*, vol. 34, no. 4, pp. 578–584, 2014.

[2] Jennifer L Dearborn, Yiyi Zhang, Ye Qiao, Muhammad Fareed K Suri, Li Liu, Rebecca F Gottesman, Andreea M Rawlings, Thomas H Mosley, Alvaro Alonso,

David S Knopman, et al., "Intracranial atherosclerosis and dementia: the atherosclerosis risk in communities (aric) study," *Neurology*, vol. 88, no. 16, pp. 1556–1563, 2017.

[3] Félix Dumais, Marco Perez Caceres, Félix Janelle, Kassem Seifeldine, Noémie Arès-Bruneau, Jose Gutierrez, Christian Bocti, and Kevin Whittingstall, "eicab: A novel deep learning pipeline for circle of willis multiclass segmentation and analysis," *Neuroimage*, vol. 260, pp. 119425, 2022.

[4] Iris N Vos, Ynte M Ruigrok, Edwin Bennink, Mireille RE Velthuis, Barbara Paic, Maud EH Ophelders, Myrthe AD Buser, Bas HM van der Velden, Geng Chen, Matthieu Coupet, et al., "Evaluation of techniques for automated classification and artery quantification of the circle of willis on tof-mra images: The crown challenge," *Medical Image Analysis*, p. 103650, 2025.

[5] Pierre Rougé, Nicolas Passat, and Odyssée Merveille, "Topology aware multitask cascaded u-net for cerebrovascular segmentation," *PloS one*, vol. 19, no. 12, pp. e0311439, 2024.

[6] Zichen Zhang, Min Tang, Dana Cobzas, Dornoosh Zonoobi, Martin Jagersand, and Jacob L Jaremko, "End-to-end detection-segmentation network with roi convolution," in *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*. IEEE, 2018, pp. 1509–1512.

[7] Kaiyuan Yang, Fabio Musio, Yihui Ma, Norman Juchler, Johannes C Paetzold, Rami Al-Maskari, Luciano Höher, Hongwei Bran Li, Ibrahim Ethem Hamamci, Anjany Sekuboyina, et al., "Benchmarking the cow with the topcow challenge: Topology-aware anatomical segmentation of the circle of willis for cta and mra," *ArXiv*, pp. arXiv–2312, 2025.

[8] Tommaso Di Noto, Guillaume Marie, Sebastien Tourbier, Yasser Aleman-Gomez, Oscar Esteban, Guillaume Saliou, Meritxell Bach Cuadra, Patric Hagmann, and Jonas Richiardi, "Towards automated brain aneurysm detection in tof-mra: open data, weak labels, and anatomical knowledge," *Neuroinformatics*, vol. 21, no. 1, pp. 21–34, 2023.

[9] Rafic Nader, Vincent L'Allinec, Romain Bourcier, and Florent Autrusseau, "Two-steps neural networks for an automated cerebrovascular landmark detection along the circle of willis," *IEEE Journal of Biomedical and Health Informatics*, 2025.

[10] Paul F Jaeger, Simon AA Kohl, Sebastian Bickelhaupt, Fabian Isensee, Tristan Anselm Kuder, Heinz-Peter Schlemmer, and Klaus H Maier-Hein, "Retina u-net: Embarrassingly simple exploitation of segmentation supervision for medical object detection," in *Machine learning for health workshop*. PMLR, 2020, pp. 171–183.