3D Mesh-based Intracranial Aneurysm Detection using a Synthetic Bifurcation Model

Said Ammari[1], Anass Nouri[1], Romain Bourcier[2], and Florent Autrusseau[2,3]

[1] SETIME Laboratory, Faculty of Sciences, Ibn Tofail University, Kenitra, Morocco
said.ammari@uit.ac.ma, anass.nouri@uit.ac.ma
[2] Institut du Thorax, CHU de Nantes, Nantes Université, Nantes, France
romain.bourcier@chu-nantes.fr
[3] LTEN, Nantes Université, Polytech'Nantes, Nantes, France
florent.autrusseau@univ-nantes.fr

Abstract. An accurate detection of Unruptured Intracranial Aneurysms (UIAs) is crucial for evaluating the risk of rupture and, possibly, scheduling appropriate treatment strategies. Time-of-Flight Magnetic Resonance Angiography (TOF-MRA) is the preferred non-invasive imaging technique for visualizing cerebral vasculatures and aneurysms. However, traditional 2D image analysis methods often fail to capture the complex 3D geometry of aneurysms. This work introduces an approach that transforms TOF-MRA images into 3D mesh representations which are subsequently analyzed via MeshCNN, a geometrical deep learning architecture designed for 3D mesh classification. To further enhance the architecture's performances, we perform data augmentation; in this step, we introduce a synthetic bifurcation model capable of realistically mimicking the shapes of vessels and aneurysms. We evaluate the effectiveness of our approach on a dataset comprising various aneurysm sizes and locations. The results demonstrated a significant improvement in detection accuracy and consistency.

Keywords: Unruptured Intracranial Aneurysms · 3D meshes · Aneurysm Detection · Deep Learning · MeshCNN · Vasculature Model

1 INTRODUCTION

Unruptured intracranial aneurysms (UIAs) are focal dilatations of the wall of intracranial arteries, most frequently located on the large vessels at the base of the brain, particularly along the Circle of Willis (CoW) [1]. Epidemiological studies estimate that about 3% of the adult population harbor an IntraCranial Aneurysm (ICA) [2]. Within the broader spectrum of cerebrovascular disease, ICAs are considered particularly serious lesions because of their potential for life-threatening complications [3]. Most UIAs remain asymptomatic and are discovered incidentally during angiographic imaging. However, rupture leads to aneurysmal subarachnoid hemorrhage [4], a severe form of stroke associated with high mortality and long-term disability. Early and reliable detection of unruptured aneurysms is therefore essential to guide, follow-up, and decide whether preventive treatment is warranted.

In current clinical practice, ICAs are mainly detected and assessed on angiographic acquisitions such as contrast-enhanced computed tomography angiography (CTA), digital subtraction angiography (DSA), 3D Rotational Angiography (3DRA), and magnetic resonance angiography (MRA), including time-of-flight MRA (TOF-MRA) for non-invasive screening. After an UIA has been identified, neuroradiologists characterize its size, neck, and relation to parent vessels using a set of geometric measurements. Although these measurements can be performed on two-dimensional projections, the underlying anatomy is inherently three-dimensional and complex, which makes manual 2D measurements sensitive to plane selection and observer variability. Three-dimensional representations of the vessel lumen therefore provide a more consistent basis for quantifying aneurysm morphology and for integrating quantitative shape indices into risk assessment.

Most recent deep-learning approaches for aneurysm detection and analysis operate directly on volumetric CTA or MRA data [5–8]. These models have shown promising performance, but they remain tightly coupled to image contrast, voxel size and acquisition protocol, and often require large modality-specific training datasets [9]. An alternative line of work consists in first segmenting the intracranial vessels and then representing the vasculature as a three-dimensional surface mesh. Vascular surface meshes are compact, encode the true three-dimensional geometry of the vessel tree, and can be derived from different imaging modalities. Deep geometric learning on such meshes offers a way to build UIAs detectors that are less dependent on image intensity and more focused on local vascular shape [10].

In this work, we build on the recent advances in geometric deep learning on vascular surface meshes and propose a mesh-based UIA detection framework specifically tailored to the CoW crops around 15 regions of interest. Our approach operates directly on vessel surface meshes extracted from 3D TOF-MRA and learns to classify local mesh regions as aneurysmal or healthy using only geometric features. Furthermore, we leverage a recent developed aneuryusmal synthetic model to significantly enhance classification performance.

We clarify the scope and the specific contribution of this work, since both MeshCNN [11] and the VaMos synthetic model [13] have been previously introduced as standalone methods. Our contribution is threefold: (i) we adapt the VaMos generative pipeline so that its synthetic binary volumes are directly transformed into surface meshes that are compatible with a geometric deep-learning architecture; (ii) we design a size-aware incremental augmentation strategy that explicitly targets the under-represented small-aneurysm regime; and (iii) we provide, on a multicentric cohort, a systematic comparison between anatomically realistic synthetic augmentation and conventional geometric perturbations (rotation, jittering).

The architecture, training strategy, and data preprocessing pipeline are described in Section II, the experimental results and comparisons to existing methods in Section III, and the implications and perspectives for clinical use in Section IV.

2 METHODS

2.1 MeshCNN

MeshCNN [11] is a convolutional neural network designed to operate directly on triangular surface meshes constituted from vertices and edges. Unlike volumetric methods that require computationally expensive voxelization, MeshCNN operates on a triangular surface mesh $\mathcal{M} = (V, E, F)$, where V, E, F denote the sets of vertices, edges, and faces, respectively. By bypassing the discretization of the 3D space, the network preserves the precise topology and high-resolution geometric details of the vascular surface. Instead of pixels or voxels, the network treats edges as basic processing units and associates to each edge $e \in E$ a local geometric descriptor (dihedral angle, two inner angles, and two edge–length ratios). These descriptors are intentionally designed to be invariant to rigid transformations—such as rotation and translation—and global scaling. This invariance is critical in medical imaging to ensure consistent detection regardless of patient orientation or variations in vessel size.

Convolutions are defined on mesh edges and operate on an edge together with its four 1-ring neighboring edges. A symmetric ordering of these neighboring edges is used to ensure that the convolution is invariant to the edge orientation and independent of the local mesh parameterization. Pooling is implemented as a learned edge-collapse procedure: edges are ranked according to the magnitude of their learned features, and the least informative ones are iteratively collapsed. This progressively simplifies the mesh while preserving its overall topology and concentrating resolution in geometrically salient regions (such as the aneurysm dome and neck). For segmentation, the corresponding unpooling layers use the recorded collapse history to up-sample features back to the original mesh resolution.

In this study, we leverage MeshCNN to analyze cerebral vessel surface meshes. We process the vascular tree using patches with a spatial extent of 19.2 mm (equivalent to 48^3 voxels at 0.4 mm resolution). This patch-based approach provides a sufficient anatomical window to capture the bifurcation geometry and the adjacent healthy parent vessels while maintaining computational efficiency.

The model follows a MeshUNet architecture, specialized for the segmentation of vascular structures. The network starts with 32 feature maps in the first layer, which are progressively doubled (64, 128, 256) through the encoder path. Each convolutional block is composed of three residual blocks to enhance feature propagation.

Geometric Feature Processing: Each 3D patch is represented by 2,330 input edges. The network utilizes three levels of learned edge-collapse pooling, progressively simplifying the mesh resolution from 2,330 edges to 1,800, 1,350, and finally 600 edges. This hierarchical simplification allows the model to capture multi-scale geometric features, effectively isolating the aneurysmal sac from the parent vessel. For the segmentation task, corresponding unpooling layers restore the features to the original 2,330-edge resolution via the recorded collapse history.

Training Protocol: The model was trained using the following hyperparameters:

- **Optimizer and Loss:** We employed the Adam optimizer with an initial learning rate of 0.001. The training objective was minimized using the Cross-Entropy loss function, which is standard for edge-wise segmentation tasks.
- **Data Augmentation:** To improve generalization and account for clinical variability, we performed 20 augmentations per sample, including a vertex sliding factor of 0.2, which introduces subtle geometric deformations without altering the mesh topology.
- **Batch and Epochs:** Training was conducted with a batch size of 12. We utilized an early stopping strategy to prevent overfitting, halting the process when the validation loss ceased to improve for a predefined number of epochs.
- **Hardware:** All experiments were performed on an NVIDIA RTX 4060 GPU (8 GB VRAM), providing the necessary memory capacity for high-dimensional mesh operations.

2.2 Synthetic aneurysm mesh model

To address the limited availability of annotated aneurysm cases, we augment our dataset with realistic synthetic samples generated using the VaMos framework [13].

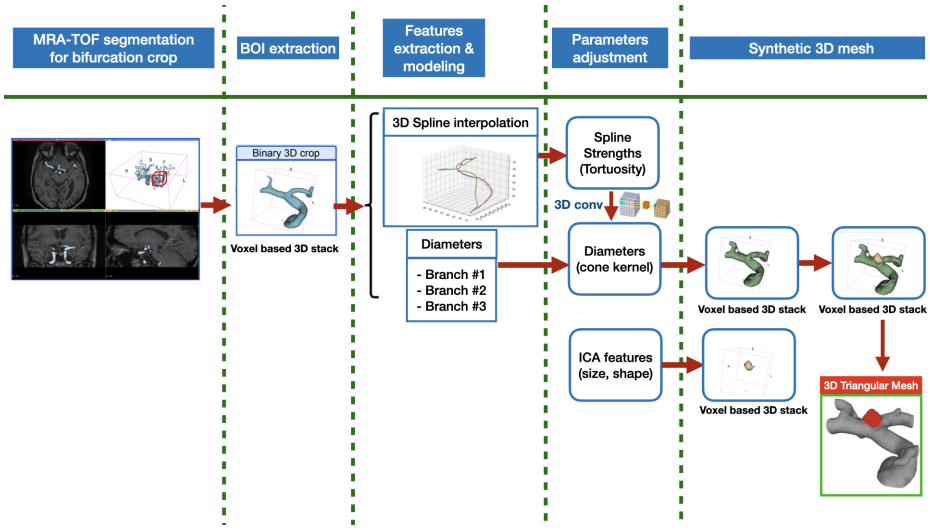

Fig. 1. Pipeline of the adapted VaMoS synthetic 3D meshes model [13]

VaMos was originally proposed as a generative model for cerebral arterial bifurcations and aneurysms derived from TOF–MRA data. Starting from a binary

vessel mask, the arterial tree is skeletonized and represented as a 3D centerline graph. Each branch of this graph is then approximated by a B-spline curve, whose control points and degree encode the local vessel geometry and tortuosity. By modifying these spline parameters, VaMos is able to generate new anatomically plausible arterial bifurcations.

In addition to geometric synthesis, VaMos can simulate realistic image textures by extracting noise characteristics from the original, unsegmented TOF-MRA volumes. As a result, the framework produces two outputs: (a) a synthetic grayscale volume containing the artery, surrounding noise, gray and dark matter, and the texture of the cerebrospinal fluid; and (b) a synthetic binary segmentation volume.

In our work, we exclusively use the generated voxel-based binary segmentation, which are subsequently converted into 3D meshes. The generation process follows a sequential pipeline illustrated in Figure 1. First, a bifurcation is cropped from a cerebral vasculature, and its 3D skeleton is computed. Each skeleton branch is then modeled using a 3D spline, allowing the spline parameters to be adjusted to modify branch tortuosity. For each branch, the convolution between the modified centerline and a generated 3D spherical kernel (including elastic distortions) allows the model to precisely control the arterial diameter and converts the skeletal representation into a volumetric structure.

The objective is to generate bifurcations that are sufficiently similar to, yet distinct from, their original counterparts, thereby promoting robust learning in the subsequent deep learning architecture.

VaMos then places a synthetic aneurysm along the bisector of the two daughters branches at a distance D from the bifurcation center. The aneurysm is modeled as a spherical structure of radius r that undergoes elastic deformation to ensure an anatomically plausible appearance.

The placement distance D is defined as [13]:

$$D = r \cdot \gamma \; + \; \sqrt{\left(\frac{R}{\tan(\theta/2)}\right)^2 + R^2}, \tag{1}$$

where *gamma* is a growth factor that controls the inward or outward displacement of the aneurysm relative to the bifurcation, R is the radius of the branches that make up the bifurcation and θ represent the angles between the daughters branches. Figure 2 illustrates the aneurysm positioning process and the parameters involved in this formulation.

Finally, the synthetic binary volumes are directly converted into triangular vessel surface meshes (e.g., using a marching cubes algorithm).

2.3 Dataset and preprocessing

Our data set consists of 190 aneurysmal bifurcations extracted from 189 TOF-MRA volumes (one patient presenting two aneurysms), together with 190 non-aneurysmal bifurcations, yielding a balanced binary classification problem. The 190 non-aneurysmal bifurcations constitute the negative class of this balanced

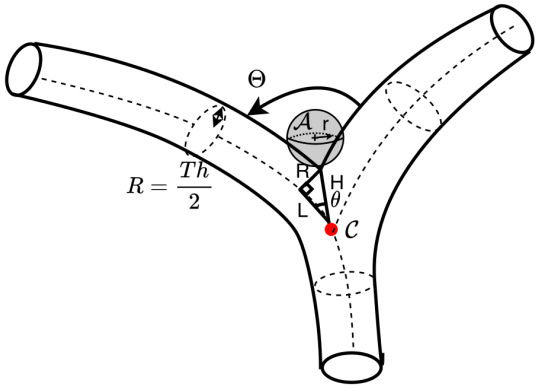

Fig. 2. Position of an ICA between the daughter branches of a synthetic bifurcation [13].

classification task: each one is a healthy arterial bifurcation cropped with the same $48 \times 48 \times 48$ voxel window as the aneurysmal cases (here centered on the bifurcation itself), and they are split across the training, validation and test sets following the same partition as the aneurysmal cases. These scans were collected from a multicentric cohort involving 19 different MRI scanners from multiple manufacturers and various French institutions. The non-aneurysmal bifurcations were extracted from healthy arterial bifurcations within the same set of TOF-MRA volumes, ensuring consistent acquisition parameters between the two classes. Using the synthetic model, we generated up to 200 additional aneurysmal bifurcations, which were distributed across different experiments in order to study the effectiveness of the synthetic data and to address the question of how much additional data is beneficial. To prevent any form of data leakage, all synthetic aneurysmal bifurcations were generated exclusively from real bifurcations belonging to the training set. A detailed description of the data set comprising the human aneurysmal bifurcation (denoted as "real" in the remaining) is provided in Table 1. The cases are well distributed across the training, validation and test sets with respect to aneurysm size. While our model leverages a synthetic bifurcation framework for training enhancement, the final assessment was performed on an independent test set consisting entirely of real clinical acquisitions. As shown in Table 1, there are only 52 small aneurysms ($R \leq 2$), which corresponds to about 27% of the real data set. Medium-sized aneurysms ($2 < R \leq 4$) represent around 62% (119 aneurysms), which is relatively high, while large aneurysms ($R > 4$) account for approximately 10% (19 aneurysms). This is not expected to be a major issue, since large aneurysms are usually easily detected by neuroradiologists. As argued in the beginning, our goal is to enrich the real database with synthetic aneurysmal bifurcations in order to optimize the learning and the inference on small aneurysms. The synthetic dataset composition is summarized in Table 2. We designed an incremental augmentation

strategy to systematically address the challenge of detecting small aneurysms, a limitation that has been widely reported in the literature [14–16].

Table 1. Number of clinical aneurysms used classified per radius category

Split	**Small** ($R \leq 2\,\text{mm}$)	**Medium** ($2 < R \leq 4\,\text{mm}$)	**Large** ($R > 4\,\text{mm}$)	**Total**
Training	25	76	11	112
Validation	9	19	2	30
Test	18	24	6	48
Total	52	119	19	190

Specifically, we conducted three experiments with progressively increasing the number of synthetic volumes: Exp2 introduced 50 synthetic aneurysms, Exp3 doubled this to 100, and Exp4 scaled the total to 200. Crucially, this generation was heavily skewed towards the 'Small' category ($R \leq 2\,\text{mm}$) to force the model's attention toward their subtle features. As shown in Table 2, the proportion of small aneurysms remains dominant across all tiers, culminating in Exp4 where 179 out of 200 (89.5%) added samples were small aneurysms. This deliberate bias aims to shift the model's optimization focus toward learning the challenging morphological signatures of small aneurysms, which are typically underrepresented in clinical datasets.

Our proprietary dataset was designed according to the following procedure. Starting from 3D TOF–MRA volumes, we first identify the bifurcations of interest. For aneurysmal bifurcations, we crop a $48 \times 48 \times 48$ voxel patch centered on the aneurysm; these crop sizes were perfectly estimated to encompass the full aneurysm, and, to some extent, the bifurcation carrying the aneurysm; for non-aneurysmal cases, the crop is centered on the bifurcation center itself. On these cropped volumes, we segment the intracranial vasculature to obtain binary vessel masks, and then extract the vessel surface using a standard marching-cubes procedure, yielding high-resolution triangular meshes of the local vascular geometry.

Table 2. Number of synthetic aneurysms added per experiment and radius category

Experiment	**Small** ($R \leq 2\,\text{mm}$)	**Medium** ($2 < R \leq 4\,\text{mm}$)	**Total**
Exp2	43	7	50
Exp3	83	17	100
Exp4	179	21	200

Before training, the meshes undergo a cleaning step to remove artifacts (e.g. duplicated internal surfaces) and to keep only the connected component corresponding to the bifurcation. The aneurysm and parent-vessel masks are then

Table 3. Detection metrics for synthetic augmentation experiments

Experiments	Sens.	Spec.	Prec.	F1
Exp1 (Real)	95.83	85.41	86.79	91.08
Exp2 (+50ICA)	100.00	85.41	87.27	93.20
Exp3 (+100ICA)	**100.00**	**89.58**	**90.56**	**95.04**
Exp4 (+200ICA)	95.83	91.66	92.00	93.87
Exp5 (+Jittering)	79.16	82.14	95.00	86.35
Exp6 (+Rotation)	91.66	45.83	62.85	74.56
Exp7 (+Jittering&Rotation)	93.75	87.5	88.23	90.90

mapped onto the mesh, providing an edge-wise ground-truth labeling (aneurysm vs. normal vessel). These labeled meshes constitute the input to MeshCNN for a supervised aneurysm detection.

3 EXPERIMENTS AND RESULTS

3.1 Impact of Synthetic Data Augmentation

We first evaluate the overall impact of integrating synthetic data into the training pipeline. To assess the potential clinical applicability of our model, we established a rigorous "use case" evaluation protocol defined by stringent Intersection over Union (IoU) thresholds. We consider a prediction to be a valid *True Positive* (detection) only if it demonstrates a high geometric overlap with the ground truth, requiring an **IoU** > 0.4. This threshold exceeds standard loose detection criteria, ensuring that only precise localizations are rewarded [17]. Conversely, we adopted an unforgiving standard for *False Positives*, particularly on aneurysm-free meshes. To strictly penalize noise and artifacts, any predicted positive region on a non-aneurysmal bifurcation exhibiting a minimal intersection area (or segmentation area) of 0.1 is immediately classified as a False Positive. By enforcing this dual-threshold strategy-demanding high fidelity for detection (> 0.4) while maintaining a low tolerance for spurious segmentation (> 0.1)- we ensure that our reported performance metrics reflect a model capable of robust detection with minimal false alarms.

Table 3 presents the performance of the baseline model (where MeshCNN is trained solely on real data) compared to models augmented with synthetic data, and to models augmented with jittering and/or rotation. As we can remark, synthetic data effectively enriches the patterns learned, providing the network with diverse geometric variations that are underrepresented in the limited real-world training set. A clear trend of improvement is observed when supplementing the real dataset with targeted synthetic samples, specifically up to an optimal saturation point:

1. **Baseline vs. Augmentation:** The baseline model (Exp1) established a strong foundation with an F1-score of 91.08% and a sensitivity of 95.83%. The introduction of synthetic positive samples (Exp2, +50ICAs) immediately enhanced model robustness. Sensitivity reached 100.00%, completely eliminating

false negatives. This suggests that even a moderate injection of well designed synthetic geometries helps the model resolve the decision boundaries.

2. **Optimal Configuration (Exp3):** The highest overall performance was achieved in Exp3 (+100ICAs). This configuration maintained perfect sensitivity (100.00%) while simultaneously improving specificity to 89.58% (+4.17% over baseline) and precision to 90.56%. Consequently, Exp3 yielded the peak of F1-score of 95.04%, demonstrating that this amount of synthetic data provides the optimal trade-off between recall and precision, effectively reducing false positives without compromising the detection of true cases.

3. **Saturation and Diminishing Returns (Exp4):** Increasing the number of synthetic data further to 200 samples (Exp4) resulted in diminishing returns. While specificity continued to rise to a maximum of 91.66%, indicating a highly precise model, sensitivity regressed to the baseline level of 95.83%. This drop caused a decrease in the F1-Score to 93.87%, suggesting that an excessive ration of synthetic data may lead to overfitting on synthetic features, potentially obscuring the subtle variations present in the real validation cases.

4. **Limitations of Traditional Augmentation:** Standard techniques like jittering (Exp5) and rotation (Exp6) performed poorly. Rotation led to a severe collapse in specificity (45.83%). Since MeshCNN's edge features are inherently invariant to rigid motions, global rotation provides no new geometric information and instead appears to hinder convergence. Similarly, stochastic vertex jittering (Exp5) corrupted the subtle surface curvatures necessary for detection, dropping sensitivity to 79.16%.

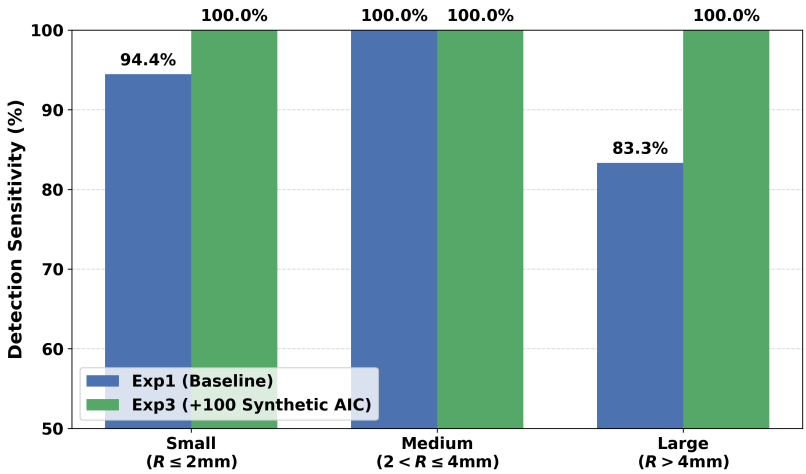

Fig. 3. Impact of Synthetic Data on Aneurysm Size Detection

3.2 Aneurysm-Size-Stratified Sensitivity Analysis

To further understand the model's performance, we analyze detection sensitivity across three clinically relevant ICAs-radii categories: Small (≤ 2 mm), Medium ($2 - 4$ mm), and Large (> 4 mm). As shown in Figure 3, the introduction of synthetic data yields the most significant performance gains in the *Large* category, with more moderate improvements observed in the *Small* category.

The baseline model already reaches a high sensitivity on small ICAs ($R \leq 2$ mm), missing only 1 out of 18 cases (94.4%); these lesions nonetheless remain the most challenging category, as they are the most prone to being confused with vessel noise or with the bifurcation itself. However, by augmenting with synthetic examples explicitly generating these challenging geometries, the model learns more robust feature descriptors, closing this residual gap and reaching a perfect sensitivity on small aneurysms. The breakdown of detection performance by aneurysm size, illustrated in Figure 3, isolates the precise contribution of synthetic augmentation. Therefore, it simultaneously removes the residual blind spot on small lesions and substantially strengthens the under-represented 'Large' category. This confirms that the synthetic data acted as a targeted corrective mechanism for the dataset's specific blind spots rather than merely providing a generic improvement in performance.

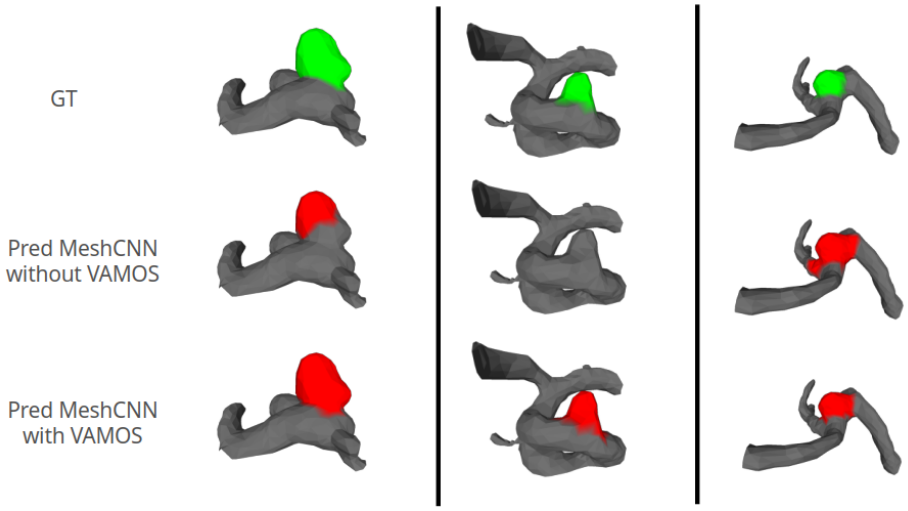

Fig. 4. Qualitative comparison of aneurysm prediction results on representative vascular mesh samples. From top to bottom: ground truth annotations (GT), predictions obtained using the baseline MeshCNN without VAMOS, and predictions obtained using MeshCNN enhanced with VAMOS. Green regions denote expert ground-truth aneurysm annotations, while red regions indicate predicted aneurysm areas.

A performance drop was observed in the detection of large intracranial aneurysms compared to smaller samples. This trend can be attributed to the inherent class imbalance within our training set, as large ICAs are significantly less prevalent in clinical databases. From a clinical perspective, large aneurysms are more easily identified by neuroradiologists experts during routine radiological screenings, leading to prompt intervention via endovascular coiling or surgical clipping. Consequently, the availability of such cases for research datasets is naturally limited.

The qualitative interest of integrating the VAMOS framework is further illustrated in Figure 4. In contrast to the baseline results (Exp1), which rely solely on real clinical data, the MeshCNN-augmented model (Exp3), thanks to VAMOS, demonstrates a superior ability to refine the detection process. While the baseline often produces unstable or fragmented predictions on complex bifurcations, the synthetic augmentation strategy significantly improves the robustness of the detection process. It specifically improves the model's capacity to resolve the spatial boundaries of the aneurysm, its neck and the bifurcation around it, ensuring that the detection remains localized and consistent even across the challenging small-radius geometries previously discussed.

4 DISCUSSION AND CONCLUSION

In this work, we introduced a mesh-based framework for the detection of ICAs from TOF-MRA, in which vascular bifurcations are represented as 3D surface meshes and processed with a geometric-based deep learning architecture (MeshCNN). In contrast to the approaches operating on grid-voxels images, our method considers solely the geometry while being complemented by a synthetic vasculature model able to generate realistic aneurysmal bifurcations. Using this synthetic model, we systematically evaluated the impact of enriching the training set with varying numbers of synthetic samples and compared these configurations to classical geometric augmentations (e.g. random jittering and rigid rotations). Our experiments demonstrate that shape-preserving synthetic augmentation leads to a significantly larger and more consistent improvement in detection performance than conventional geometric perturbations. Notably, while rotation and vertex jittering are staples of 2D computer vision, they proved counterproductive in this context. We observed that rigid rotations provided no additional informative value due to the inherent rotation-invariance of MeshCNN's input descriptors, while stochastic jittering tended to corrupt the subtle surface curvatures essential for identifying the aneurysm neck and dome. By explicitly enriching the learning phase with synthetic cases, the MeshCNN model becomes more sensitive to subtle aneurysmal deformations. In its current form, the proposed pipeline operates on pre-localized regions of interest cropped around candidate bifurcations and therefore addresses a patch-level aneurysm classification task rather than a fully automated end-to-end detection problem on complete TOF-MRA volumes. Perspectives include integrating a vessel localization stage into the processing pipeline, evaluating the framework on larger and more hetero-

geneous cohorts using systematic cross-validation, and validating the proposed synthetic augmentation strategy on Computed Tomography Angiography (CTA) scans to assess cross-modality generalization. In addition, we plan to refine the modeling and supervision of aneurysm neck geometry and to couple detection with more detailed morphometric characterization (e.g. radius, aspect ratio, non-sphericity), with the aim of better supporting early detection of small ICAs, reducing false negative detections, and improving downstream clinical decision-making.

5 ACKNOWLEDGMENTS

This work was supported by the National Center for Scientific and Technical Research (CNRST) of Morocco through the PhD-Associate Scholarship.

References

1. R. D. Brown and J. P. Broderick, "Unruptured intracranial aneurysms: epidemiology, natural history, management options, and familial screening," *The Lancet Neurology*, vol. 13, no. 4, pp. 393–404, 2014.
2. M. H. M. Vlak, A. Algra, R. Brandenburg and G. J. E. Rinkel, "Prevalence of unruptured intracranial aneurysms, with emphasis on sex, age, comorbidity, country, and time period: A systematic review and meta-analysis," *Lancet Neurology*, vol. 10, pp. 626–636, 2011.
3. R. Pascalau, V. A. Padurean, D. Bartos, A. Bartos, and B. A. Szabo, "The geometry of the circle of Willis anatomical variants as a potential cerebrovascular risk factor," *Turk Neurosurg*, vol. 29, no. 2,pp. 151–158, 2019.
4. D. J. Nieuwkamp, L. E. Setz, A. Algra, F. H. H. Linn, N. K. de Rooij and G. J. E. Rinkel, "Changes in case fatality of aneurysmal subarachnoid haemorrhage over time, according to age, sex, and region: A meta-analysis," *Lancet Neurology*, vol. 8, pp. 635–642, 2009.
5. A. Nouri, F. Autrusseau, R. Bourcier, A. Gaignard, V. L'allinec, C. Menguy, J. Véziers, H. Desal, G. Loirand, and R. Redon, "Characterization of 3D bifurcations in micro-scan and MRA-TOF images of cerebral vasculature for prediction of intra-cranial aneurysms," *Computerized Medical Imaging and Graphics*, vol. 84, p. 101751, 2020.
6. B. Joo, S. S. Ahn, P. H. Yoon, S. Bae, B. Sohn, Y. E. Lee, J. H. Bae, M. S. Park, H. S. Choi and S. K. Lee, "A deep learning algorithm may automate intracranial aneurysm detection on MR angiography with high diagnostic performance," *European Radiology*, vol. 30, pp. 5785–5793, 2020.
7. T. Sichtermann, A. Faron, R. Sijben, N. Teichert, J. Freiherr and M. Wiesmann, "Deep learning–based detection of intracranial aneurysms in 3D TOF-MRA," *American Journal of Neuroradiology*, vol. 40, pp. 25–32, 2019.
8. R. Nader, R. Bourcier, and F. Autrusseau, "Using deep learning for an automatic detection and classification of the vascular bifurcations along the Circle of Willis," *Medical Image Analysis*, vol. 89, p. 102919, 2023.
9. G. Litjens, T. Kooi, B. E. Bejnordi *et al.*, "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.

10. K. M. Timmins *et al.*, "Deep Learning with Vessel Surface Meshes for Intracranial Aneurysm Detection," in *Proc. SPIE Medical Imaging*, 2022.
11. Hanocka, Rana, et al. "Meshcnn: a network with an edge." ACM Transactions on Graphics (ToG) 38.4 (2019): 1-12.
12. F. Autrusseau, R. Nader, A. Nouri, V. L'Allinec and R. Bourcier, "Toward a 3D arterial tree bifurcation model for intra-cranial aneurysm detection and segmentation," in *Proc. 26th Int. Conf. on Pattern Recognition (ICPR)*, 2022, pp. 4500–4506.
13. R. Nader, F. Autrusseau, V. L'Allinec, and R. Bourcier, "Building a Synthetic Vascular Model: Evaluation in an Intracranial Aneurysms Detection Scenario," *IEEE Transactions on Medical Imaging*, vol. 44, no. 3, pp. 1347–1358, 2025.
14. K. M. Timmins *et al.*, "Geometric Deep Learning Using Vascular Surface Meshes for Modality-Independent Unruptured Intracranial Aneurysm Detection," *IEEE Trans. Med. Imaging*, 2023.
15. Z. Zhou, Y. Jin, H. Ye, X. Zhang, J. Liu, and W. Zhang, "Classification, detection, and segmentation performance of image-based AI in intracranial aneurysm: a systematic review," *BMC Medical Imaging*, vol. 24, no. 1, p. 164, 2024.
16. I. L. Štepán-Buksakowska, J. M. Accurso, F. E. Diehn, J. Huston, T. J. Kaufmann, P. H. Luetmer, C. P. Wood, X. Yang, D. J. Blezek, R. Carter, C. Hagen, D. Hořínek, A. Hejčl, M. Roček, and B. J. Erickson, "Computer-aided diagnosis improves detection of small intracranial aneurysms on MRA in a clinical setting," *American Journal of Neuroradiology*, vol. 35, no. 10, pp. 1897–1902, 2014.
17. Ceballos-Arroyo, A.M., et al.: Anatomically-guided masked autoencoder pretraining for aneurysm detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2026).