Exploiting Interior Geometry for Intracranial Aneurysm Detection and Segmentation on 3D Point Clouds via GNNs

Mohamed Amine Layachi[1,*], Anass Nouri[1], Souad Eddarouich[2], and Florent Autrusseau[3]

[1] Ibn Tofail University, Kenitra, Morocco
[2] CRMEF, Rabat, Morocco
[3] École Polytechnique, University of Nantes, Nantes, France
mohamedamine.layachi@uit.ac.ma, anass.nouri@uit.ac.ma, eddarouichsouad@gmail.com, florent.autrusseau@univ-nantes.fr

Abstract. The automated 3D detection and segmentation of intracranial aneurysms is critical for clinical intervention, yet existing deep learning methods often struggle with the complex, tubular organization of cerebral vasculature. We propose EXPIGEO, a geometry-aware framework that bridges this gap by integrating explicit geometric reasoning with topological graph representation. Standard local features are easily confounded by the inherent curvilinearity of the vasculature, heavily limiting the effectiveness of conventional Graph Neural Networks (GNNs). EXPIGEO overcomes this by introducing an internal geometric exploration mechanism to characterize the vessel interior, enabling the robust extraction of skeletal topology and physiological descriptors. This explicit structural encoding successfully resolves topological ambiguities and unlocks GNN performance, establishing a new state-of-the-art benchmark for aneurysm detection on the IntrA dataset with 98.41% accuracy and an F1-score of 0.9589. Beyond raw performance, EXPIGEO provides clinical transparency: instance-level explainability confirms that predictions are driven by biologically relevant volumetric features, ensuring a reliable and interpretable diagnosis.

Keywords: Intracranial Aneurysm · Detection and Segmentation · 3D Point Clouds · Graph Neural Networks · Explainability

1 Introduction

An Intracranial Aneurysm (IA) is a pathological dilation of cerebral arteries [1]. Morphologically, IAs are primarily classified into *saccular* (berry) aneurysms—focal outpouchings at arterial bifurcations [2]—and rare *fusiform* aneurysms that circumferentially involve the entire vessel wall segment [3, 4]. While often asymptomatic, rupture precipitates subarachnoid hemorrhage (SAH), a catastrophic event with high morbidity [5]. Consequently, precise 3D delineation of

* Corresponding author: mohamedamine.layachi@uit.ac.ma

these complex features is paramount for rupture risk assessment. In this context, point cloud–based learning has emerged as a compelling alternative to voxel-driven pipelines, offering high geometric fidelity for sparse vascular structures with significantly reduced computational overhead.

Several approaches tackle IA detection and segmentation directly on point clouds. Early efforts [6] employ PointNet [7] architectures to learn permutation-invariant representations across varying modalities. Addressing annotation scarcity, [8] propose an unsupervised contrastive framework using dual-branch encoders [9]. For precise segmentation, [10] adopt a cascaded strategy, classifying fragments before fine-grained delineation via SO-Net [11]. Recent works explore multi-modal fusion, combining spatial geometry with 2D appearance cues [12], or introduce task-specific inductive biases, such as attention mechanisms for long-range context [13] and dual-stream transformers targeting boundary precision [14].

Despite their natural fit for topological data, Graph Neural Networks remain largely underexplored in vascular analysis. Existing point-based GNN architectures [15, 16] are often optimized for urban environments and rely heavily on auxiliary convolutional layers. Their limited medical adoption stems from two primary challenges. First, standard local descriptors (e.g., spatial coordinates, surface curvature) struggle to distinguish pathological dilations from healthy vessels due to the inherent, complex curvilinearity of the vascular network. Consequently, simplistic node features necessitate complex hybrid models. Furthermore, deep GNNs suffer from structural pathologies: *over-smoothing*, where node representations lose distinctiveness, and *over-squashing*, which results in the loss of long-range dependencies. Once geometric ambiguities are resolved, GNNs' explicit topological structure enables post-hoc explainability. This offers a critical advantage over opaque black-box models in clinical settings where transparency and trustworthiness are paramount.

To bridge this gap, we propose EXPIGEO[4], a novel geometry-aware framework for point-based 3D intracranial aneurysm analysis. Our pipeline overcomes the limitations of standard local features by operating in four stages: (1) Interior Geometry Exploration to robustly extract the vascular skeleton; (2) Geometric Descriptors Extraction to compute high-fidelity, physiology-aware features; (3) Graph Construction and Architecture, where we build a feature-rich graph to train a GNN for dense node-level classification, from which lesion detection is inherently derived; and (4) Explainability, applying instance-level methods to isolate decision-driving subgraphs. To our knowledge, EXPIGEO is the first framework to successfully unlock GNN performance for 3D aneurysm analysis while integrating post-hoc explainability, providing a transparent lens into the network's predictive behavior. The code, a companion video on interior geometry exploration, and further visualizations are available at: https://github.com/mohamedaminelayachi/EXPIGEO.

[4] EXPIGEO: EXPloitable Interior GEOmetry.

2 Methodology

As illustrated in Figure 1, the proposed framework operates as an automated, end-to-end pipeline. First, we employ *Ray Probing* to explore the point cloud's interior. Second, to overcome the vulnerability of standard surface-based algorithms to complex tubular geometry, we extract a robust centerline to guide the derivation of the vascular skeleton and compute structural descriptors. Third, to deliver the exact 3D boundaries required for clinical surgical planning, the network operates directly on the full surface point cloud. By structuring this surface as a feature-rich graph, a GNN performs binary point-wise classification (segmentation), from which instance-level aneurysm identification (detection) is subsequently derived. Finally, we apply instance-level explainability to quantify feature importance and diagnose the structural origins of classification errors within the model's receptive field.

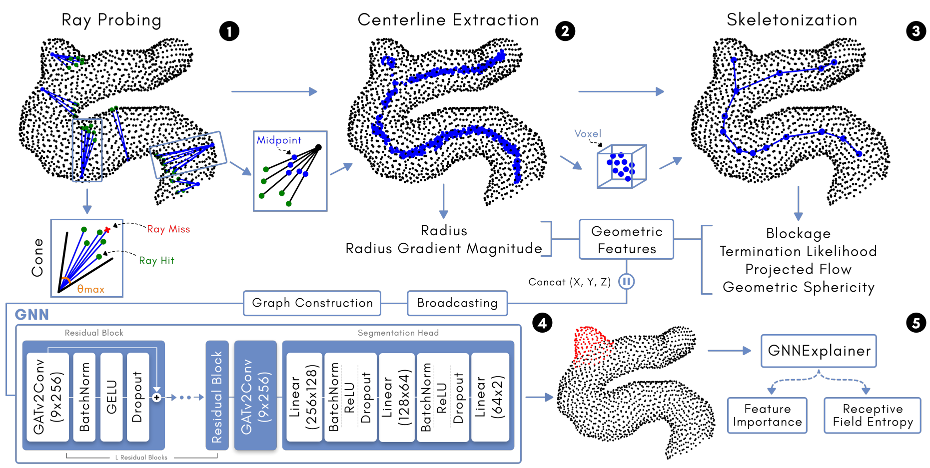

Fig. 1. Overview of the EXPIGEO Framework.

2.1 Interior Geometry Exploration

Ray Probing. We utilize the surface point cloud $\mathcal{P}_{\text{surface}} = \{(\boldsymbol{p}_i, \boldsymbol{n}_i)\}_{i=1}^{N}$, comprising point coordinates $\boldsymbol{p}_i \in \mathbb{R}^3$ and surface normal vectors $\boldsymbol{n}_i \in \mathbb{S}^2$. From a source $(\boldsymbol{o}, \boldsymbol{n}_o) \in \mathcal{P}_{\text{surface}}$, we define an inward cone of directions $\boldsymbol{d} \in \mathbb{S}^2$ satisfying $\boldsymbol{d}^\top(-\boldsymbol{n}_o) \geq \cos(\theta_{\max})$. A ray $\boldsymbol{r}(s) = \boldsymbol{o} + s\boldsymbol{d}$ (where $s > 0$) is considered to *hit* a target $\boldsymbol{t} \in \{\boldsymbol{p}_i\}$ if the perpendicular distance $\|\boldsymbol{t} - (\boldsymbol{o} + s^*\boldsymbol{d})\| \leq \delta$, where $s^* = \boldsymbol{d}^\top(\boldsymbol{t} - \boldsymbol{o})$ is the projection of the target onto the ray. The threshold is defined as $\delta = \gamma \cdot \bar{d}_k$, where $\bar{d}_k$ is the mean Euclidean distance to the k-nearest neighbors and γ is a tolerance factor. We employ KD-Trees for spatial queries to bypass naive $\mathcal{O}(N^2)$ *hit* testing, achieving $\mathcal{O}(N \log N)$ time and $\mathcal{O}(N)$ space complexity. We sample $K = 8$ rays per cone with $\theta_{\max} = 15°$ and $\gamma = 1/3$.

Centerline Extraction & Radius Estimation. Let $\mathcal{H}_i$ denote the set of K rays cast from a source point $\boldsymbol{o}_i \in \mathcal{P}_{\text{surface}}$. We identify the valid subset of finite intersections $\mathcal{H}'_i \subset \mathbb{R}^3$ to estimate the centerline point cloud $\mathcal{P}_{\text{center}} = \{\boldsymbol{c}_i \in \mathbb{R}^3\}$ and radius field $\mathcal{R} = \{r_i \in \mathbb{R}\}$. Computed directly from the intersection points, the estimates are defined as:

$$\boldsymbol{c}_i = \frac{\boldsymbol{o}_i + \tilde{\boldsymbol{t}}_i}{2}, \quad r_i = \frac{1}{2} \operatorname*{median}_{\boldsymbol{t} \in \mathcal{H}'_i} \|\boldsymbol{t} - \boldsymbol{o}_i\| \ , \tag{1}$$

where $\tilde{\boldsymbol{t}}_i$ is the spatial median of $\mathcal{H}'_i$. To ensure geometric fidelity, we retain a candidate $\boldsymbol{c}_i, r_i$ only if it satisfies hit coherence $\sigma_i < \tau_\sigma$ and spatial density $\bar{d}_i < \mu_d + \eta\sigma_d$. Here, σ_i denotes the hit dispersion (standard deviation of distances in $\mathcal{H}'_i$), and $\bar{d}_i$ is the local spacing computed over k neighbors. The parameters μ_d and σ_d represent the mean and standard deviation of the spacing distribution $\{\bar{d}_i\}$ across the entire cloud, while $\eta > 0$ is a tolerance factor. In our experiments, we set $k = 10$ and $\eta = 0.05$.

Skeletonization. While 3D volumetric thinning [17] is a foundational approach, extracting skeletons directly from the surface of tubular point clouds is highly sensitive to boundary noise due to the complex, overlapping geometry of vasculatures, which inevitably generates erroneous spurious branches. To mitigate this, we bypass the noisy surface and derive a robust skeleton directly from the extracted centerline $\mathcal{P}_{\text{center}}$ by voxelizing the cloud into a set of vertices $V = \{\boldsymbol{v}_j\}_{j=1}^{M}$, representing the spatial median of voxels with density ρ_j. We construct a Minimum Spanning Tree on V and iteratively prune leaf vertices satisfying incident edge length $l_j < \tau_l$ and density $\rho_j < \tau_\rho$, which act as thresholds to filter out short spurious branches and low-confidence outliers, respectively. Finally, proximal vertices are merged by spatial averaging using an adaptive threshold τ_{merge} that scales the local spacing by λ (set to 0.7):

$$\tau_{\text{merge}} = \lambda \cdot \operatorname*{median}_{\boldsymbol{v} \in V} \left(\min_{\boldsymbol{u} \neq \boldsymbol{v}} \|\boldsymbol{v} - \boldsymbol{u}\| \right) \ . \tag{2}$$

2.2 Geometric Descriptors Extraction

Radius Gradient Magnitude. Aneurysms are characterized by abrupt radius changes along the parent vessel. We quantify this by computing the gradient magnitude of the radius field $\mathcal{R}$ using Weighted Least Squares (WLS). For each centerline point $\boldsymbol{c}_i \in \mathcal{P}_{\text{center}}$ with radius r_i, we identify the k nearest neighbors $\mathcal{N}_i$ and seek the gradient vector $\boldsymbol{g}_i \in \mathbb{R}^3$ that minimizes:

$$J(\boldsymbol{g}_i) = \sum_{j \in \mathcal{N}_i} w_j \left[(\boldsymbol{c}_j - \boldsymbol{c}_i)^\top \boldsymbol{g}_i - (r_j - r_i) \right]^2 \ . \tag{3}$$

The weights follow a Gaussian kernel $w_j = \exp(-\|\boldsymbol{c}_i - \boldsymbol{c}_j\|^2 / 2\sigma_k^2)$ to discount distant points, where the scale $\sigma_k = \text{mean}_{j \in \mathcal{N}_i} \|\boldsymbol{c}_j - \boldsymbol{c}_i\|$ is the average distance to the neighbors. Solving this linear system for every $\boldsymbol{c}_i$ yields the gradient field, from which we extract the pointwise magnitudes $\|\nabla \mathcal{R}\| = \{\|\boldsymbol{g}_i\|\}$.

Termination Likelihood & Blockage. To characterize vessel termination and distinguish structural closures (e.g., aneurysms) from functional outlets, we employ two complementary descriptors. First, we compute a *Termination Likelihood* $P_{\text{term}}(\boldsymbol{v})$ for every skeletal vertex $\boldsymbol{v} \in V$ to softly highlight terminal regions. Modeled as a Gaussian function of the geodesic distance d_G to the nearest leaf vertex set $L \subset V$, the probability is defined as:

$$P_{\text{term}}(\boldsymbol{v}) = \exp\left(-\frac{\min_{\boldsymbol{u}\in L} d_G(\boldsymbol{v}, \boldsymbol{u})^2}{2\sigma^2}\right) \quad , \tag{4}$$

where σ controls the spatial fall-off of the probability field (set to 2.0). Second, to differentiate physical caps from open cut-offs, we calculate a binary *Blockage* Score $\beta_{\boldsymbol{v}}$ exclusively at leaf nodes. For a leaf $\boldsymbol{v}$ with predecessor $\boldsymbol{u}$, we define the flow direction $\boldsymbol{d}_{\boldsymbol{v}} = (\boldsymbol{v} - \boldsymbol{u})/\|\boldsymbol{v} - \boldsymbol{u}\|$ and inspect the local surface cluster $\Omega_{\boldsymbol{v}}$. The vertex $\boldsymbol{v}$ is classified as obstructed ($\beta_{\boldsymbol{v}} = 1$) if any neighbor $\boldsymbol{q} \in \Omega_{\boldsymbol{v}}$ falls within a forward aperture cone, satisfying $(\boldsymbol{q} - \boldsymbol{v})^\top \boldsymbol{d}_{\boldsymbol{v}}/\|\boldsymbol{q} - \boldsymbol{v}\| \geq \cos(\theta)$.

Geometric Sphericity. To capture local geometry independent of skeletal topology, we quantify the resemblance of the entire vascular surface to spherical versus cylindrical models. For each skeletal vertex $\boldsymbol{v}$ across the network and its mapped surface neighborhood $\Omega_{\boldsymbol{v}}$, we compute the set of centered vectors $\{\boldsymbol{q}_i = \boldsymbol{p}_i - \boldsymbol{v} \mid \boldsymbol{p}_i \in \Omega_{\boldsymbol{v}}\}$. We define the *Sphere Error* as the standard deviation of the vector norms, $\epsilon_{\text{sph}} = \text{std}(\{\|\boldsymbol{q}_i\|\})$. Simultaneously, we perform PCA on the cluster to extract principal axes $\{\boldsymbol{u}_1, \boldsymbol{u}_2, \boldsymbol{u}_3\}$. The *Cylinder Error* is computed as the minimum standard deviation of the radial distances to the candidate axes ($\boldsymbol{u}_1$ or $\boldsymbol{u}_3$): $\epsilon_{\text{cyl}} = \min_{\boldsymbol{a}\in\{\boldsymbol{u}_1,\boldsymbol{u}_3\}} \text{std}(\{\|\boldsymbol{q}_i \times \boldsymbol{a}\|\})$. The *Geometric Sphericity* descriptor is computed as the ratio:

$$\phi(\boldsymbol{v}) = \frac{\epsilon_{\text{cyl}}}{\epsilon_{\text{cyl}} + \epsilon_{\text{sph}} + \epsilon} \quad , \tag{5}$$

where ϵ is a stability constant. This metric efficiently highlights spherical structures (low ϵ_{sph}) typical of aneurysms.

Projected Flow. To contextualize local surface geometry relative to the vessel's path, we compute the *Projected Flow* $F_{\text{proj}}(\boldsymbol{p}) = (\boldsymbol{p} - \boldsymbol{v}_p)^\top \boldsymbol{d}_{\boldsymbol{v}_p}$. This value represents the scalar projection of the relative position of a surface point $\boldsymbol{p}$ onto the skeletal direction $\boldsymbol{d}_{\boldsymbol{v}_p}$ of its nearest vertex $\boldsymbol{v}_p$, effectively encoding the point's longitudinal alignment along the flow axis.

2.3 Graph Construction and Architecture

Graph Construction. We construct a graph where each point $\boldsymbol{p}_i \in \mathcal{P}_{\text{surface}}$ constitutes a node. To bridge the dense surface and sparse skeleton V, a nearest-neighbor mapping $\psi(\boldsymbol{p}_i) = \text{argmin}_{\boldsymbol{v}\in V} \|\boldsymbol{p}_i - \boldsymbol{v}\|$ induces a Voronoi partition to broadcast descriptors. Computed across the entire input bifurcation prior to downstream segmentation and detection tasks, these attributes populate a feature matrix $X \in \mathbb{R}^{N\times 9}$. Each node's feature vector concatenates its coordinates and broadcast descriptors as $\boldsymbol{x}_i = [\boldsymbol{p}_i, P_{\text{term}}, \beta, \phi, r, \|\nabla\mathcal{R}\|, F_{\text{proj}}]^\top$. Finally, k-NN links each node to its $k = 32$ closest Euclidean neighbors in $\mathcal{P}_{\text{surface}}$.

Architecture. To mitigate the risk of over-smoothing, we adopt a minimalist architecture inspired by [18]. As shown in Figure 1, our backbone comprises L residual blocks, each containing a single GATv2Conv layer [19] configured with two attention heads to overcome the static attention limitations of standard GATConv [20]. Within each block, the graph convolution is sequentially followed by Batch Normalization, a GELU activation, and Dropout. The resulting node embeddings are then processed by a three-layer Multi-Layer Perceptron (MLP) segmentation head. Each MLP layer, excluding the final output, is followed by Batch Normalization, a ReLU activation, and Dropout. This MLP progressively halves the feature dimension before mapping to the final binary predictions (aneurysm vs. vessel).

2.4 Explainability

To analyze the model's decision-making, we employ GNNExplainer [21], a technique that identifies critical input features and the relevant subgraph driving predictions via optimization. Since the explanation is inherently constrained by the message-passing scope, we examine the k-hop receptive field $\Gamma_k(\boldsymbol{p})$ around a target point $\boldsymbol{p}$ as an upper bound for the explanatory subgraph. To quantify topological ambiguity within this region, we compute the *Receptive Field Entropy*:

$$H_{\text{RF}}(\boldsymbol{p}) = -\sum_{c=0}^{C-1} P(c)\log_2 P(c), \quad \text{where } P(c) = \frac{1}{|\Gamma_k(\boldsymbol{p})|}\sum_{\boldsymbol{q}\in\Gamma_k(\boldsymbol{p})} \mathbb{1}(y_{\boldsymbol{q}} = c)\ . \tag{6}$$

Here, $P(c)$ represents the probability mass of class c within the receptive field. A high entropy indicates that $\boldsymbol{p}$ is situated at a decision boundary where message passing aggregates conflicting class signals.

3 Experiments

3.1 Dataset, Evaluation Metrics, and Implementation Details

We evaluate our framework on the IntrA dataset [22], processing localized bifurcations (aneurysmal vs. non-aneurysmal) rather than the entire vasculature. Using 116 annotated aneurysmal bifurcations and healthy counterparts, a 4-fold cross-validation partition yields 175 training (86 positives, 89 negatives), 23 validation (13 positives, 10 negatives), and 34 testing (17 positives, 17 negatives) samples. Input features are Z-score normalized. Implemented in PyTorch Geometric (NVIDIA Tesla T4), the model ($L = 4$ blocks, $\sim$719k parameters) is trained for 40 epochs ($\approx$25 mins, batch size 4) using AdamW (LR

Table 1. Quantitative comparison of segmentation performance.

Methods	Vessel		Aneurysm		mIoU (%)	mDSC (%)
	IoU (%)	DSC (%)	IoU (%)	DSC (%)		
PointNet++	93.42	96.48	76.22	83.92	84.82	90.20
SO-Net	94.46	97.09	81.40	88.76	87.93	92.92
PointConv	94.65	97.18	79.53	86.52	87.09	91.85
3DMedPT	94.82	97.29	81.80	89.25	88.31	93.27
PT	95.10	97.49	82.11	90.18	88.61	93.84
EPT-Net	95.46	97.68	**87.58**	**93.38**	**91.52**	**95.53**
EXPIGEO (Ours)	**97.711**	**98.81**	85.26	91.40	91.48	95.10
Coord. Only	66.48	77.84	25.13	35.72	45.80	56.78
Coord. and Curvature	69.36	79.90	26.54	36.24	47.95	58.07

Table 2. Quantitative comparison of detection performance.

Methods	Vessel Accuracy (%)	Aneurysm Accuracy (%)	F1-Score
PointCNN	98.95	85.81	0.9044
PointNet++	98.52	88.51	0.9029
SO-Net	98.76	84.24	0.8840
Unsupervised Dual Branch	95.41	89.47	0.8226
EXPIGEO (Ours)	**99.86**	**98.41**	**0.9589**

10^{-4}, WD 10^{-5}). For binary node classification, we minimize $\mathcal{L} = \mathrm{WCE}(y, \hat{y}) + \alpha\mathcal{L}_{\mathrm{Lovasz}}(y, \hat{y})$, with $\alpha = 0.7$ balancing Weighted Cross-Entropy (WCE) and Lovász-Softmax [23] losses. These predictions drive two tasks: (1) Segmentation, assessed via point-level Intersection over Union (IoU) and the Dice Similarity Coefficient (DSC); and (2) Detection, where candidates require ≥10% IoU to be valid True Positives, evaluated via Accuracy, Sensitivity, Specificity, Precision, and F1-Score.

3.2 Results

Quantitative Analysis. Tables 1 and 2 show our efficacy against state-of-the-art methods [9, 11, 13, 14, 8, 24–26]. EXPIGEO achieves state-of-the-art vessel labeling (97.71%±0.31% IoU) and competitive aneurysm IoU (85.26%±3.41%), trailing only the parameter-heavy EPT-Net while remaining highly efficient: average inference is 0.227 ± 0.058s (CPU) for ray-probing per 512-point cloud and 13.1 ± 2.8ms for the GNN (GPU). Ablation proves our descriptors' necessity: a baseline GNN using only coordinates yields 25.13% on aneurysms. Adding Gaussian curvature offers negligible improvement (26.54%), as curvilinear vasculature confounds standard local features. EXPIGEO's structural encoding drives a 58.7% absolute performance leap, establishing a new detection benchmark with 98.41% accuracy—outperforming the next best method by ∼9%—and a 0.9589 F1-Score. Clinically, the model delivers 92.86% sensitivity and 99.06% specificity, ensuring reliable identification with minimized false alarms (93.55% precision). Figure 2 illustrates these results.

Clinical Insights. Structural analysis reveals that high-entropy receptive fields exacerbate over-smoothing and directly correlate with performance deterioration. Specifically, 90% of false positives exhibited an elevated mean entropy ($\bar{H}_{RF} = 0.788$), confirming that ambiguous topological neighborhoods precipitate misclassification. To dissect the drivers behind these decisions, we computed global feature importance scores using GNNExplainer. The results validate our explicit geometric modeling: while geometric sphericity (ϕ) remains universally vital for both classes (scoring $\approx$0.93) due to inherent vessel curvilinearity, distinct class-specific dependencies emerge. Most notably, radius (r) and blockage (β) are significantly more critical for aneurysm detection (scoring 1.0 and 0.55, respectively) than for healthy vessels (0.93 and 0.26), perfectly reflecting the physiological reality of volumetric expansion and flow containment.

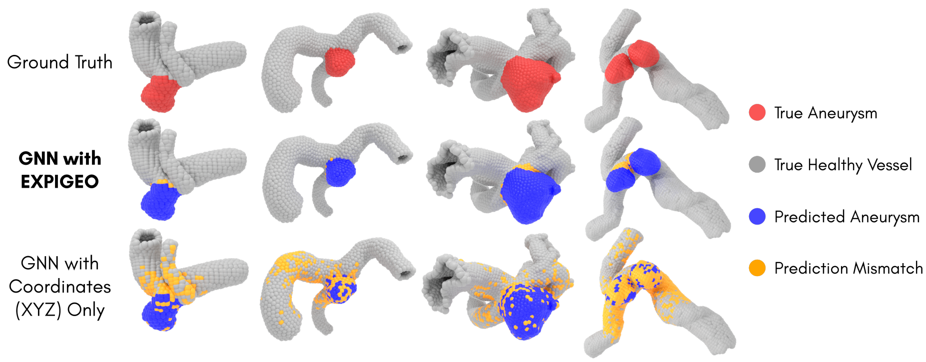

Fig. 2. Illustrative examples of EXPIGEO compared with ground truth and a spatial coordinates-only baseline GNN.

4 Conclusion

In this work, we presented EXPIGEO, a geometry-aware framework for the detection and segmentation of intracranial aneurysms on 3D point clouds. By explicitly modeling vessel topology through ray probing and skeletal feature extraction, we overcome the inherent topological ambiguities that confound standard GNNs. This structural encoding enables our architecture to establish a new state-of-the-art in aneurysm detection while maintaining clinical transparency. Through instance-level explainability, we demonstrated that the model's predictions are driven by physiologically relevant features rather than spurious correlations. Future work will focus on designing more advanced geometric descriptors and developing specialized GNN architectures tailored specifically for this downstream vascular task.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Schievink, Wouter I.: Intracranial Aneurysms. In: New England Journal of Medicine. pp. 28-40 (1997). https://doi.org/10.1056/nejm199701023360106
2. Keedy A: An overview of intracranial aneurysms. In: McGill Journal of Medicine (2020). https://doi.org/10.26443/mjm.v9i2.672
3. Park SH, Yim MB, Lee CY, Kim E, Son EI: Intracranial Fusiform Aneurysms: It's Pathogenesis, Clinical Characteristics and Managements. In: Journal of Korean Neurosurgical Society. pp. 116-123 (2008). https://doi.org/10.3340/jkns.2008.44.3.116
4. Bonneville F, Sourour N, Biondi A: Intracranial Aneurysms: an Overview. In: Neuroimaging Clinics of North America. pp. 371-382 (2006). https://doi.org/10.1016/j.nic.2006.05.001
5. D'Souza, Stanlies: Aneurysmal Subarachnoid Hemorrhage. In: Journal of Neurosurgical Anesthesiology. pp 222-240 (2015). https://doi.org/10.1097/ana.0000000000000130
6. Žiga Bizjak and Boštjan Likar and Franjo Pernuš and Žiga Špiclin: Modality agnostic intracranial aneurysm detection through supervised vascular surface classification. In: Medical Imaging (2021). https://doi.org/10.1117/12.2580868
7. Qi, Charles R and Su, Hao and Mo, Kaichun and Guibas, Leonidas J: Pointnet: Deep Learning on Point Sets for 3D Classification and Segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 77-85 (2017). https://doi.org/10.1109/CVPR.2017.16
8. Shao, Di and Lu, Xuequan and Liu, Xiao: 3D Intracranial Aneurysm Classification and Segmentation via Unsupervised Dual-Branch Learning. In: IEEE Journal of Biomedical and Health Informatics. pp. 1770-1779 (2023). https://doi.org/10.1109/JBHI.2022.3180326
9. Qi, Charles R. and Yi, Li and Su, Hao and Guibas, Leonidas J.: PointNet++: deep hierarchical feature learning on point sets in a metric space. In: Proceedings of the International Conference on Neural Information Processing Systems. pp. 5105–5114 (2017). https://api.semanticscholar.org/CorpusID:1745976
10. Yang, Xi and Xia, Ding and Kin, Taichi and Igarashi, Takeo: A two-step surface-based 3D deep learning pipeline for segmentation of intracranial aneurysms. In: Computational Visual Media. pp. 57–69 (2023). https://doi.org/10.1007/s41095-022-0270-z
11. Li, Jiaxin and Chen, Ben M. and Lee, Gim Hee: SO-Net: Self-Organizing Network for Point Cloud Analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9397-9406 (2018). https://doi.org/10.1109/CVPR.2018.00979
12. Cao, Ruifen and Zhang, Dongwei and Wei, Pijing and Ding, Yun and Zheng, Chunhou and Tan, Dayu and Zhou, Chao: PMMNet: A Dual Branch Fusion Network of Point Cloud and Multi-View for Intracranial Aneurysm Classification and Segmentation. In: IEEE Journal of Biomedical and Health Informatics. pp. 3137-3147 (2025). https://doi.org/10.1109/JBHI.2024.3380054
13. Yu, Jianhui and Zhang, Chaoyi and Wang, Heng and Zhang, Dingxin and Song, Yang and Xiang, Tiange and Liu, Dongnan and Cai, Weidong: 3D Medical Point Transformer: Introducing Convolution to Attention Networks for Medical Point Cloud Analysis (2021). ArXiv. https://api.semanticscholar.org/CorpusID:245006079
14. Liu, Yifan and Liu, Jie and Yuan, Yixuan: Edge-Oriented Point-Cloud Transformer for 3D intracranial aneurysm segmentation. In: Medical Image Computing and Computer Assisted Intervention - MICCAI 2022. Lecture Notes in Computer Science,

vol 13435. Springer, Cham. pp. 97–106 (2022). https://doi.org/10.1007/978-3-031-16443-9_10

15. Shi, Weijing and Rajkumar, Ragunathan (Raj): Point-GNN: Graph Neural Network for 3D Object Detection in a Point Cloud. In: Proceedings of the The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1708-1716 (2020). https://doi.org/10.1109/CVPR42600.2020.00178
16. Lin, Zhi-Hao and Huang, Sheng-Yu and Wang, Yu-Chiang Frank: Convolution in the Cloud: Learning Deformable Kernels in 3D Graph Convolution Networks for Point Cloud Analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1797-1806 (2020). https://doi.org/10.1109/CVPR42600.2020.00187
17. T.C. Lee and R.L. Kashyap and C.N. Chu: Building Skeleton Models via 3-D Medial Surface Axis Thinning Algorithms. In: CVGIP: Graphical Models and Image Processing. pp. 462-478 (1994). https://doi.org/10.1006/cgip.1994.1042
18. Martin, H. and Bucher, D. and Hong, Y. and Buffat, R. and Rupprecht, C. and Raubal, M. and Escalante, H. J.a and Hadsell, R.: Graph-ResNets for short-term traffic forecasts in almost unknown cities. In: Proceedings of Machine Learning Research. pp. 153–163 (2020)
19. Shaked, Brody and Uri, Alon and Eran, Yahav: How Attentive are Graph Attention Networks?. International Conference on Learning Representations (ICLR) (2022). https://api.semanticscholar.org/CorpusID:235254358
20. Veličković, Petar and Cucurull, Guillem and Casanova, Arantxa and Romero, Adriana and Liò, Pietro and Bengio, Yoshua: Graph Attention Networks. International Conference on Learning Representations (ICLR) (2018). https://doi.org/10.17863/CAM.48429
21. Ying, Rex and Bourgeois, Dylan and You, Jiaxuan and Zitnik, Marinka and Leskovec, Jure: GNNExplainer: Generating Explanations for Graph Neural Networks. Conference on Neural Information Processing Systems (NeurIPS). pp. 9240-9251 (2019). https://api.semanticscholar.org/CorpusID:202572927
22. Yang, Xi and Xia, Ding and Kin, Taichi and Igarashi, Takeo: IntrA: 3D Intracranial Aneurysm Dataset for Deep Learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2653-2663 (2020). https://doi.org/10.1109/CVPR42600.2020.00273
23. Berman, Maxim and Rannen Triki, Amal and Blaschko, Matthew B: The Lovász-Softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4413-4421 (2018). https://doi.org/10.1109/CVPR.2018.00464
24. Wu, Wenxuan and Qi, Zhongang and Fuxin, Li: PointConv: Deep Convolutional Networks on 3D Point Clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9613-9622 (2019). https://doi.org/10.1109/CVPR.2019.00985
25. Zhao, Hengshuang and Jiang, Li and Jia, Jiaya and Torr, Philip and Koltun, Vladlen: Point Transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision (ICCV). pp. 16239-16248 (2021). https://doi.org/10.1109/ICCV48922.2021.01595
26. Li, Yangyan and Bu, Rui and Sun, Mingchao and Wu, Wei and Di, Xinhan and Chen, Baoquan: PointCNN: Convolution On $\mathcal{X}$-Transformed Points. In: Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS). pp. 828–838 (2018). https://api.semanticscholar.org/CorpusID:53399839