Towards Reliable Intracranial Aneurysm Analysis: Leveraging Foundation Models for Structure-Aware Detection and Segmentation across 3D CTA, ToF-MRA and DSA

by

Muzammal Shafique
muzshf@umich.edu
https://orcid.org/0009-0000-8048-5746

A dissertation submitted in fulfillment
of the requirements for the degree of
Masters in Computer Science
The University of Michigan
2026

MS Thesis Committee:

Professor Khalid M. Malik, Chair
Dr. Makram Soui
Dr. Mohammad-Reza Siadat (External Evaluator)

DEDICATION

In the name of Allah, the Most Gracious, the Most Merciful,

To Prophet Muhammad PBUH,
whose life is a timeless source of guidance and mercy for mankind.

To my parents,

whose sacrifices, prayers, and unwavering belief in me
made every step of this journey possible.

And to the millions living with intracranial aneurysms,
may this work, however small, bring us closer
to a future where no life is lost to a silent rupture.

ACKNOWLEDGEMENTS

First and foremost, I express my deepest gratitude to my thesis advisor and committee chair, **Professor Dr. Khalid M. Malik**, for his exceptional mentorship, intellectual generosity, and steadfast support throughout this research. His vision for clinically meaningful AI and his trust in my ideas shaped every aspect of this work. I am truly fortunate to have had a supervisor who challenged me to think deeper and always pushed me toward rigor and impact.

I extend my sincere thanks to my committee members, **Dr. Makram Soui** and **Dr. Mohammad-Reza Siadat** as an external evaluator, for their valuable feedback, thoughtful guidance, and constructive engagement during my thesis defense and beyond. Their diverse expertise enriched this work in ways I deeply appreciate.

I owe a special debt of gratitude to **Dr. Ghaus Malik, M.D.** of Department of Neurosurgery at Henry Ford Health System, for generously sharing his clinical expertise and time. His deep knowledge of aneurysm morphology, rupture biomarkers, cerebral artery anatomy, and surgical decision-making provided invaluable grounding for the computational methods developed in this thesis. Collaborating with him reinforced why this research matters.

I am grateful to the University of Michigan and the Great Lakes HPC cluster for providing the computational resources that made these experiments possible.

Finally, to my family: thank you for your patience, love, and encouragement. Every late night and early morning was sustained by your support.

PREFACE

This thesis is submitted in partial fulfillment of the requirements for the degree of Master of Science in Computer Science at the University of Michigan-Flint. The research presented here was conducted under the supervision of **Professor Dr. Khalid M. Malik** at the Scalable Medical Imaging and Learning Systems Laboratory **(SMILES Lab)**, in collaboration with Henry Ford Health System.

The work addresses a problem I encountered early in my graduate studies: despite the life-threatening nature of intracranial aneurysms, the computational tools available to clinicians for automated analysis remained limited - either ignoring the anatomical context of the vascular system, or failing to delineate aneurysm boundaries with the precision required for rupture risk assessment. This motivated the development of a clinically grounded, two-component framework: **ARAN**, for structure-aware aneurysm detection and artery-of-origin identification from non-invasive computed tomography angiography (CTA) and time-of-flight magnetic resonance angiography (ToF-MRA) scans, and **FocusSDF**, for boundary-precise aneurysm segmentation from high-resolution DSA images.

Parts of this thesis have been accepted by the top-tier peer-reviewed venues for publication. All DSA modality data used in this research was handled in accordance with institutional review board (IRB) protocols, and in-house patient data from Henry Ford Health System was used under IRB No. 11254. Publicly available datasets of CTA and ToF-MRA were also used to conduct experiments.

This work sits at the intersection of medical image analysis, geometric deep learning, and clinical neurovascular imaging. It is my hope that the methods developed here serve not only as technical contributions, but as meaningful steps toward AI systems that clinicians can trust and patients can benefit from.

Muzammal Shafique
University of Michigan
April 23, 2026

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF ACRONYMS

IA Intracranial Aneurysm

CTA Computed Tomography Angiography

DSA Digital Subtraction Angiography

ToF-MRA Time-of-Flight Magnetic Resonance Angiography

MRI Magnetic Resonance Imaging

ARAN **AR**tery-aware intracranial **AN**eurysm detection

ICA Internal Carotid Artery

ACom Anterior Communicating Artery

MCA Middle Cerebral Artery

PCom Posterior Communicating Artery

BA Basilar Artery

PCA Posterior Cerebral Artery

ROI Region of Interest

GGCA Geometry-Gated Cross-Attention

LIST OF SYMBOLS

$\mathcal{X}$	Input 3D angiographic scan volume.
H, W, D	Spatial dimensions (height, width, depth) of the input volume.
f_θ	Encoder network of the foundation model.
$\mathcal{Z}$	Latent feature representation produced by the encoder.
C	Number of feature channels in the latent representation.
H', W', D'	Spatial dimensions of the encoded feature map.
$g(\cdot)$	Decoder (segmentation head) network.
$\hat{Y}$	Predicted voxel-wise artery class probabilities.
c_k	k-th point on the vessel centerline.
A_k	Cross-sectional area at centerline point c_k.
ϱ_k	Estimated vessel radius at point c_k.
$\mathcal{R}$	Radius profile along the vessel centerline.
α_k	Major axis of fitted ellipse at cross-section k.
ξ_k	Minor axis of fitted ellipse at cross-section k.
ϵ_k	Cross-sectional eccentricity at point k.
$\mathbf{c}(s)$	Parametric representation of vessel centerline curve.
s	Arc-length parameter along the centerline.
$\mathbf{T}(s)$	Tangent vector of the curve.
$\mathbf{N}(s)$	Normal vector of the curve.
$\mathbf{B}(s)$	Binormal vector of the curve.
$\kappa(s)$	Curvature of the vessel centerline.

$\tau(s)$	Torsion of the vessel centerline.
$\mathcal{G}$	Vascular graph representation.
$\mathcal{V}$	Set of nodes (arteries) in the graph.
$\mathcal{E}$	Set of edges representing artery relationships.
u_k	Node corresponding to the k-th artery.
$\mathbf{h}_k$	Feature vector associated with node k.
$\bar{\varrho}_k$	Mean vessel radius for artery k.
$\bar{\epsilon}_k$	Mean eccentricity for artery k.
$\bar{\kappa}_k$	Mean curvature for artery k.
$\bar{\tau}_k$	Mean torsion for artery k.
$\mathcal{N}(k)$	Neighborhood of node k in the graph.
$\mathbf{W}^{(l)}$	Learnable weight matrix at layer l in the graph network.
σ	Nonlinear activation function.
$\mathbf{f}_{\text{vis}}$	Volumetric visual feature map from encoder.
Λ	Voxel grid of the image volume.
$\mathcal{R}$	Region of interest (ROI) for a specific artery.
$\mathbf{p}$	Voxel location within the ROI.
$\mathbf{z}_{\text{vis}}$	ROI-pooled visual feature vector.
$\mathbf{W}_v$	Learnable projection matrix for visual features.
$\mathbf{e}_{\text{vis}}$	Visual embedding.
$\mathbf{e}_{\text{geo}}$	Geometry-aware embedding.
d	Embedding dimension.
$\mathbf{Q}, \mathbf{K}, \mathbf{V}$	Query, key, and value vectors in cross-attention.
$\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V$	Learnable projection matrices for attention.
$\mathbf{a}$	Cross-attended feature representation.
$\mathbf{g}$	Geometry-conditioned gating vector.
$\mathbf{W}_g$	Learnable gating projection matrix.

$\mathbf{e}_{\text{out}}$	Fused output embedding.
$\odot$	Element-wise multiplication.
ϕ_f	Feed-forward network.
$\mathrm{LayerNorm}$	Layer normalization operation.
$\mathcal{D}$	Spatial domain (image region).
$\partial\mathcal{D}$	Boundary of the region.
$\mathbf{q}$	Spatial point in the domain.
$\mathbf{b}$	Boundary point.
$\rho(\mathbf{q}, \partial\mathcal{D})$	Euclidean distance from point $\mathbf{q}$ to boundary.
$\psi(\mathbf{q})$	Signed distance function (SDF) value at point $\mathbf{q}$.
$\mathcal{F}$	Ground truth SDF map.
$\hat{\mathcal{F}}$	Predicted SDF map.
$\mathcal{F}_j, \hat{\mathcal{F}}_j$	SDF values at pixel j.
$\mathcal{D}_{\text{in}}$	Interior region of the object.
$\mathcal{D}_{\text{out}}$	Exterior region of the object.
v_j	Spatial weight assigned to pixel j.
∇	Gradient operator.
$\mathcal{L}_{\text{FocusSDF}}$	Proposed boundary-aware loss function.
p	Norm order (L_1 or L_2).
β	Weight decay parameter for distance-based weighting.
μ	Balancing coefficient for gradient consistency term.

ABSTRACT

Intracranial aneurysm (IA) analysis is a critical task in neurovascular imaging due to its direct role in preventing subarachnoid hemorrhage and related life-threatening complications. Accurate characterization of aneurysms requires not only precise morphological delineation but also reliable anatomical localization within the cerebral vasculature, as rupture risk is strongly influenced by both aneurysm geometry and its parent artery context. Clinical imaging modalities such as Digital Subtraction Angiography (DSA), Computed Tomography Angiography (CTA), and Time-of-Flight Magnetic Resonance Angiography (ToF-MRA) provide complementary information for aneurysm assessment. While DSA remains the gold standard for high-resolution vascular visualization, it is invasive and limited in routine screening. In contrast, CTA and ToF-MRA offer non-invasive alternatives but introduce challenges such as lower spatial resolution, complex vascular overlap, and ambiguous boundary definitions, making automated interpretation difficult. Recent advances in deep learning and foundation models have significantly improved medical image understanding; however, two fundamental challenges remain insufficiently addressed: (i) insufficient incorporation of vascular structural context for clinically meaningful aneurysm localization, and (ii) lack of explicit boundary modeling for accurate lesion segmentation. This thesis addresses these limitations through a unified framework that integrates structure-aware and geometry-aware learning mechanisms for intracranial aneurysm analysis across multi-modal imaging data.

First, to address the challenge of structure-agnostic aneurysm detection in non-invasive angiographic imaging, we propose **ARAN**: **AR**tery-aware intracranial **AN**eurysm detection, a framework designed to explicitly model vascular topology and aneurysm-artery relationships. Unlike prior methods that formulate aneurysm detection as a binary classification task, ARAN incorporates arterial context as a first-class representation for improved clinical interpretability and risk assessment. The framework leverages a fine-tuned VISTA3D foundation model for volumetric feature extraction from CTA and ToF-MRA, combined with a graph-based vascular modeling branch that encodes vessel centerline geometry. This branch captures clinically relevant geometric descriptors, including vessel radius profiles, eccentricity, and Frenet-Serret curvature-torsion sequences, enabling fine-grained modeling of vascular structure. These geometric representations are fused with deep visual features using a Geometry-Gated Cross-Attention mechanism (GGCA), allowing structural priors to modulate learned image representations in a context-aware manner. Experimental results on publicly available datasets demonstrate that ARAN achieves state-of-

the-art artery-level classification performance, significantly outperforming both vision-only and geometry-only baselines, with accuracies of 78.2% on CTA and 85.8% on ToF-MRA. These results highlight the importance of integrating vascular structure modeling with foundation model representations for robust aneurysm localization.

Second, once an aneurysm is detected and its parent artery identified, precise morphological delineation of the aneurysm sac becomes essential for rupture risk assessment. To this end, we propose **FocusSDF**, a boundary-aware segmentation framework that leverages signed distance function (SDF)-based supervision to explicitly encode geometric proximity to lesion boundaries. Unlike conventional segmentation approaches that treat all pixels equally or rely on region-based supervision, FocusSDF adaptively reweights learning based on spatial distance to the anatomical boundary, thereby directing the model's attention toward structurally informative regions. This formulation improves boundary sharpness and segmentation fidelity across diverse imaging modalities, including aneurysm DSA, liver CTA, cerebral stroke MRI, and breast tumor ultrasound datasets. Extensive evaluations demonstrate that FocusSDF consistently improves performance over state-of-the-art segmentation architectures, including transformer-based and foundation model-driven approaches such as MedSAM. Together, ARAN and FocusSDF form a clinically coherent pipeline: from non-invasive aneurysm detection and artery-of-origin identification, to high-fidelity segmentation from DSA, laying the groundwork for artery-aware rupture risk prediction as a promising direction for future work.

CHAPTER 1

Introduction

1.1 Background and Motivation

Intracranial aneurysms (IA) are localized dilations of cerebral blood vessels that arise due to structural weakening of the arterial wall (1). Globally, intracranial aneurysms represent a significant health burden, causing nearly 500,000 deaths each year (2), with approximately half of the victims younger than 50, and an estimated 180 million individuals living with unruptured aneurysms worldwide (3). These abnormalities are clinically significant because their rupture can lead to subarachnoid hemorrhage, a catastrophic neurological event associated with high mortality and long-term disability. Early detection and accurate characterization of aneurysms are therefore essential for timely clinical intervention and risk assessment. In modern neurovascular practice, imaging plays a central role in both screening and diagnostic workflows, enabling non-invasive visualization of cerebrovascular structures and pathological formations.

Among available imaging modalities, CTA and ToF-MRA have become increasingly important in clinical practice due to their non-invasive nature (4). These modalities enable rapid, safe, and widely accessible visualization of cerebrovascular structures and are commonly used for initial aneurysm detection and follow-up monitoring. However, CTA and ToF-MRA suffer from several inherent limitations, including lower spatial resolution, imaging artifacts, partial volume effects, and complex overlapping vascular structures (5). These factors make automated interpretation particularly challenging, especially for small aneurysms or those located at complex vascular bifurcations. Once an aneurysm is detected through non-invasive screening, patients are typically referred for DSA, which is widely regarded as the clinical gold standard for intracranial vascular assessment. DSA provides detailed high-resolution visualization of vessel lumen and aneurysm morphology, making it highly effective for evaluating fine structural boundaries and hemodynamic properties essential for rupture risk estimation. However, DSA is invasive, requiring arterial catheterization, and is therefore associated with procedural risks such as stroke, bleeding, and patient discomfort, making it unsuitable for routine population-scale screening.

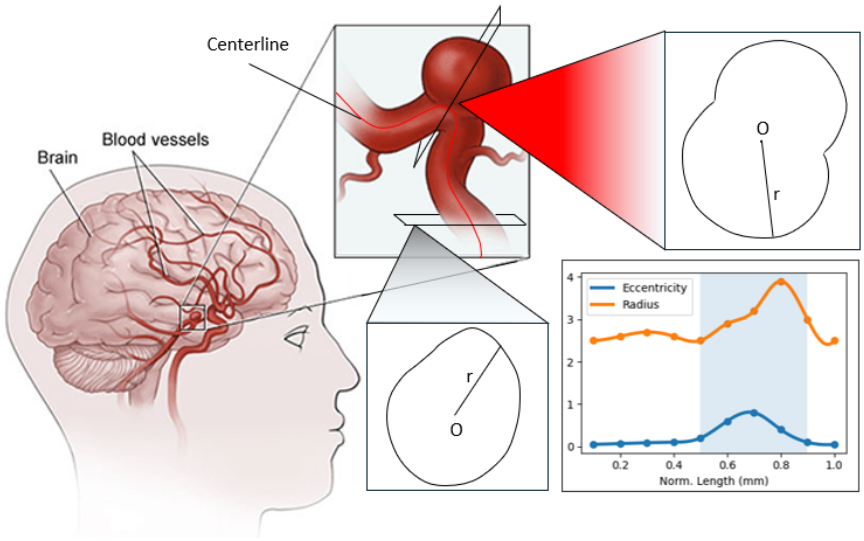

Figure 1.1: Illustration of vessel centerline (O) extraction and geometry-based feature computation. Radius (r) and cross-sectional eccentricity (e) are measured along the artery, with notable variations observed near aneurysm regions, reflecting local vessel dilation.

The increasing availability of large-scale medical imaging datasets has driven the adoption of deep learning-based methods for medical image analysis. Convolutional neural networks and, more recently, transformer-based architectures have demonstrated strong performance in tasks such as segmentation, detection, and classification across multiple medical domains. In parallel, foundation models have emerged as powerful general-purpose representations trained on large and diverse datasets, enabling improved transferability and robustness in downstream tasks. Foundation models such as SAM (6), MedSAM (7) and volumetric encoders like VISTA3D (8) have shown promising results in adapting to medical imaging tasks with limited labeled data.

Despite these advancements, two fundamental challenges remain insufficiently addressed in the context of intracranial aneurysm analysis. First, many detection-based approaches in CTA and ToF-MRA treat aneurysm identification as an isolated classification problem, without explicitly modeling the underlying vascular anatomy (9; 10). However, aneurysm behavior and rupture risk are strongly dependent on their anatomical location within the cerebral arterial tree, particularly their parent vessel and local vascular geometry (11). Reliable identification of the artery of origin is therefore a clinically essential step, as it directly informs surgical planning and risk stratification. Second, once an aneurysm is localized and a DSA scan is acquired, most existing segmentation methods do not explicitly incorporate boundary-sensitive learning mechanisms, resulting in suboptimal delineation of lesion boundaries. Accurate boundary modeling is particularly critical in aneurysm analysis, where subtle morphological variations can significantly influence rupture risk estimation and treatment decisions.

This motivates the need for a framework that integrates both global vascular structural reasoning, as illustrated in Figure 1.1, and fine-grained geometric boundary understanding, mirroring the natural clinical workflow: first detecting and localizing the aneurysm within the cerebrovascular network from non-invasive scans, and subsequently segmenting it with high boundary precision from DSA for downstream risk assessment. From a clinical perspective, such a framework should not only be capable of contextualizing aneurysm findings within the surrounding vascular network, but also of accurately delineating aneurysm boundaries to support decision-making in diagnosis, surgical planning, and rupture risk prediction (12).

1.2 Problem Statement

Despite significant advances in medical image analysis using deep learning and foundation models, intracranial aneurysm analysis remains a fundamentally challenging problem due to the complex interplay between global vascular topology and fine-grained geometric boundary structures. Existing approaches typically address either detection or segmentation in isolation, and rarely incorporate clinically meaningful anatomical constraints that are critical for accurate diagnosis and risk assessment.

First, existing aneurysm detection frameworks largely formulate the problem as a binary classification or object detection task, treating aneurysms as isolated entities without explicitly modeling their anatomical context within the cerebrovascular tree. However, clinical studies have consistently shown that aneurysm behavior and rupture probability are highly dependent on their parent artery, branching patterns, and local hemodynamic environment. Despite this, most computational approaches fail to incorporate vascular structural properties such as radius profiles, curvature, and torsion explicitly, leading to limited interpretability and reduced clinical utility. This gap becomes even more pronounced in non-invasive imaging modalities such as CTA and ToF-MRA, where vascular overlap and complex 3D geometry further complicate anatomically consistent localization. Second, once an aneurysm is detected and a high-resolution DSA scan is acquired for detailed morphological assessment, precise boundary delineation of the aneurysm sac becomes essential for rupture risk estimation. As illustrated in Fig. 1.2, medical image segmentation across different modalities poses distinct and significant challenges: in DSA, aneurysms are frequently obscured behind cluttered arterial loops with similar dye contrasts (Fig. 1.2a); in CTA, overlapping soft tissues and low tissue contrast make organ boundary delineation ambiguous (Fig. 1.2b); in MRI, stroke lesions exhibit nearly identical contrast to surrounding healthy tissue with highly irregular boundaries (Fig. 1.2c); and in ultrasound, breast tumors are structurally similar to healthy tissue, making precise delineation difficult (Fig. 1.2d). Yet most state-of-the-art segmentation methods - including convolutional neural networks, transformer-based architectures, and even recent foun-

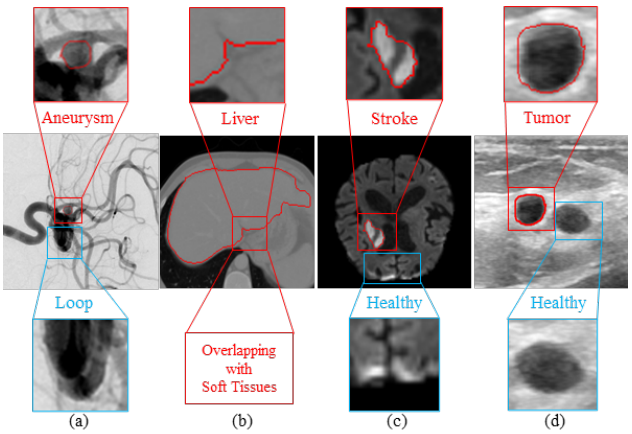

Figure 1.2: Challenges in medical image segmentation across different modalities: (a) DSA: a partially hidden aneurysm behind the cluttered arteries and loops having similar dye contrasts, (b) CTA: overlapping soft tissues with liver due to low tissue contrast, (c) MRI: irregular boundaries and identical contrast of stroke and healthy tissue, and (d) Ultrasound: similar tumor and healthy tissue structure in breast.

dation models such as MedSAM - primarily rely on region-based supervision. These methods optimize pixel-wise similarity or region overlap without explicitly modeling boundary uncertainty, resulting in blurred or imprecise segmentation boundaries. This limitation is especially critical in intracranial aneurysms, where subtle boundary variations can significantly alter morphological measurements such as dome height, neck width, and aspect ratio, which are directly linked to rupture risk assessment.

Third, while foundation models have demonstrated strong generalization capabilities across medical imaging tasks, their application to neurovascular analysis remains limited by the absence of domain-specific inductive biases. These models are typically trained on large-scale generic datasets and lack explicit mechanisms to encode vascular structural priors or geometric boundary constraints. Consequently, their predictions may not fully align with clinically relevant anatomical structures unless guided by task-specific learning strategies tailored to the neurovascular domain.

Therefore, there exists a critical need for a unified computational framework that addresses the following fundamental challenges:

1. **Structure-aware vascular reasoning:** Incorporating the topology and geometry of cerebral vessels to enable artery-level understanding and clinically meaningful localization of aneurysms from non-invasive CTA and ToF-MRA scans.

2. **Boundary-aware geometric modeling:** Enabling precise delineation of aneurysm and lesion boundaries across diverse imaging modalities, particularly in low-contrast and ambiguous regions as highlighted in Fig. 1.2.

3. **Foundation model adaptation for domain specificity:** Effectively integrating pre-trained foundation model representations with medical domain priors to improve robustness and generalization across neurovascular tasks.

4. **Cross-modal generalization:** Ensuring consistent performance across heterogeneous imaging modalities such as CTA, ToF-MRA, and DSA, each with distinct noise characteristics, resolution properties, and clinical roles.

Addressing these challenges is essential for developing clinically reliable AI systems that move beyond isolated detection or segmentation tasks and instead provide anatomically and functionally meaningful interpretations of intracranial aneurysms across the full clinical imaging pipeline. This thesis addresses these limitations through a unified framework that integrates structure-aware modeling and geometry-aware learning, both built upon foundation model backbones for improved generalization and robustness.

1.3 Significance of the Study

The significance of this research lies in its contribution to both clinical neurovascular imaging and the broader field of medical artificial intelligence, particularly in the integration of foundation models with domain-specific structural and geometric priors across the full aneurysm analysis pipeline.

From a clinical perspective, intracranial aneurysm rupture remains a major cause of subarachnoid hemorrhage, which is associated with high mortality rates and severe long-term neurological impairment among survivors. Early detection, accurate anatomical localization, and precise morphological characterization of aneurysms are therefore critical for preventive intervention and informed treatment planning. However, current clinical workflows rely heavily on expert radiologists to interpret complex imaging data across multiple modalities, which is both time-consuming and subject to inter-observer variability. By developing automated, reliable, and anatomically informed AI systems, this research has the potential to significantly assist clinicians in improving diagnostic accuracy, reducing workload, and enabling earlier risk stratification.

A key clinical contribution of this work is the introduction of artery-aware aneurysm localization from non-invasive CTA and ToF-MRA scans, which represents a significant advancement over traditional detection-based methods that treat aneurysms as isolated entities. By explicitly incorporating vascular topology and parent artery relationships through graph-based structural modeling, the proposed framework enables more meaningful anatomical interpretation of aneurysm location. This is particularly important for surgical planning and endovascular treatment strategies, where the choice of intervention depends heavily on the aneurysm's relationship to surrounding vessels

and its artery of origin. Furthermore, identifying the parent artery at the detection stage directly supports downstream risk stratification, as rupture probability is strongly correlated with the specific arterial segment involved.

Building upon this localization capability, a second key clinical contribution of this work is the improvement of boundary-level precision in aneurysm segmentation from high-resolution DSA scans, which directly impacts quantitative measurements used in clinical decision-making. More accurate boundary delineation enables better estimation of morphological features such as aneurysm size, shape irregularity, and neck-to-dome ratio, all of which are established predictors of rupture risk. By explicitly modeling boundary information through signed distance function-based supervision, the proposed approach enhances the reliability of segmentation outputs in clinically sensitive regions, particularly where aneurysm margins are low-contrast, partially obscured, or geometrically complex.

From a technological standpoint, this thesis contributes to the advancement of foundation model adaptation in medical imaging. While foundation models have demonstrated impressive generalization capabilities, their effectiveness in specialized domains such as neurovascular imaging depends critically on the integration of domain-specific inductive biases. This work demonstrates that augmenting foundation model representations with vascular structural priors for detection, and geometric boundary constraints for segmentation, significantly enhances both performance and interpretability in complex neurovascular scenarios.

Furthermore, the proposed framework supports cross-modal generalization across CTA, ToF-MRA, and DSA, enabling robust performance across both non-invasive screening modalities and the invasive gold-standard, reflecting the natural clinical imaging pathway. This is particularly important for real-world clinical deployment, where imaging conditions vary significantly across institutions and acquisition protocols.

From a broader AI research perspective, this study contributes to the emerging paradigm of physically and anatomically informed deep learning, where models are not solely data-driven but are guided by meaningful structural and geometric constraints. The integration of graph-based vascular modeling (13) for structure-aware detection and boundary-aware signed distance supervision for geometry-aware segmentation represents a step toward more interpretable and clinically aligned AI systems.

1.4 Research Objectives

The primary objective of this thesis is to develop a unified framework for intracranial aneurysm analysis that integrates structure-aware modeling and geometry-aware learning within a foundation model-driven paradigm, as illustrated in Fig. 1.3. The proposed framework is built upon a clinically

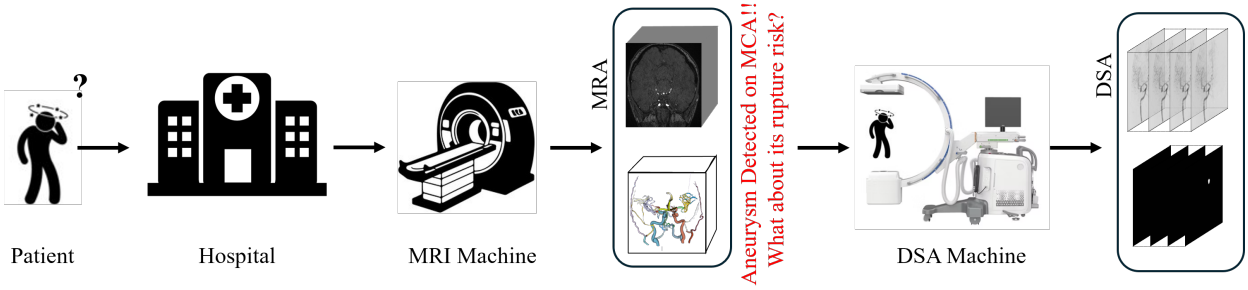

Figure 1.3: Clinical motivation and research objectives of this thesis. A patient presenting with neurological symptoms undergoes non-invasive CTA or ToF-MRA imaging, from which an aneurysm is detected and its artery of origin is identified. Given the detected aneurysm and the associated clinical concern regarding rupture risk, the patient is referred for a high-resolution DSA scan, from which the aneurysm sac is precisely segmented.

grounded pipeline that mirrors real-world neurovascular practice: first, an aneurysm is detected and its artery of origin is identified from non-invasive CTA or ToF-MRA scans; subsequently, once rupture risk concern is raised, the patient is referred for a high-resolution DSA scan from which the aneurysm sac is precisely segmented to support morphological assessment and risk quantification. By explicitly modeling the aneurysm-artery relationship and integrating boundary-sensitive geometric supervision, this thesis aims to move beyond isolated lesion detection toward structurally informed and clinically meaningful aneurysm characterization. Specifically, this work addresses key limitations in existing methods related to vascular structural reasoning, boundary modeling, and cross-modal generalization across CTA, ToF-MRA, and DSA imaging modalities. To achieve this overarching goal, the following specific research objectives are defined:

1. **To develop an artery-aware aneurysm detection framework that incorporates vascular structural priors:**
   This objective focuses on modeling aneurysm-artery relationships by integrating vascular geometry and topology from non-invasive CTA and ToF-MRA scans, enabling the identification of parent artery associations that are critical for clinical interpretation, surgical planning, and rupture risk stratification.

2. **To develop a boundary-aware loss function for precise aneurysm segmentation:**
   A key yet frequently overlooked factor in achieving this precision is the design of the loss function: model performance depends not only on the architecture but equally on the loss formulation - suboptimal losses can constrain even powerful models, whereas well-designed losses can significantly enhance simpler ones (14). This objective therefore focuses on improving the accuracy of aneurysm boundary delineation in high-resolution DSA scans by

explicitly incorporating geometric boundary information through signed distance function-based supervision.

3. **To integrate foundation model representations for improved generalization and robustness:**
   This includes adapting pre-trained medical foundation models for both 3D angiographic detection and 2D boundary-aware segmentation, ensuring that learned representations generalize across heterogeneous datasets and diverse imaging modalities with limited labeled data.

4. **To evaluate the proposed methods across multi-modal neurovascular imaging datasets:**
   This objective aims to systematically validate the proposed approaches on publicly available and in-house CTA, ToF-MRA, and DSA datasets, ensuring robustness, clinical relevance, and generalization capability across varied acquisition protocols and institutional settings.

1.5 Research Questions

The following research questions are formulated to guide the methodology and experimental design of this thesis. Each question is directly linked to the proposed contributions and is evaluated in the experimental chapters.

1. **RQ1: Structure-Aware Aneurysm Detection**: Can explicitly modeling vascular structure and parent artery relationships improve the accuracy and clinical interpretability of intracranial aneurysm detection in CTA and ToF-MRA?

2. **RQ2: Boundary-Aware Segmentation**: How can boundary-aware supervision using signed distance functions improve the accuracy and sharpness of intracranial aneurysm and medical image segmentation compared to conventional region-based and distance-transform-based loss functions?

3. **RQ3: Foundation Model Adaptation**: Do foundation model-based representations (2D and 3D) improve generalization performance in intracranial aneurysm analysis when combined with domain-specific geometric and structural priors?

1.6 Thesis Organization

This thesis is organized into six major chapters, each addressing a specific component of the proposed research framework.

1. **Chapter 1: Introduction**: This chapter introduces the clinical and technical background of intracranial aneurysm analysis, highlights existing challenges in detection and segmentation, and presents the motivation, problem statement, research objectives, and research questions. It also outlines the overall structure of the thesis.

2. **Chapter 2: Literature Review**: This chapter provides a comprehensive review of existing work in medical image segmentation, aneurysm detection, vascular modeling, and foundation models in medical imaging.

3. **Chapter 3: Proposed Methodology**: This chapter presents the unified architecture of the proposed system, integrating structure-aware detection (ARAN) and geometry-aware segmentation (FocusSDF). It describes the overall system design, data flow, and interaction between 2D and 3D components within a foundation model-driven pipeline.

4. **Chapter 4: Datasets**: This chapter describes the datasets used in this study across different imaging modalities, including CTA, ToF-MRA, and DSA. It details dataset characteristics and the selection criteria for artery segmentation, artery-aware aneurysm detection, and boundary-aware segmentation tasks.

5. **Chapter 5: Experiments and Results**: This chapter describes the datasets, evaluation metrics, experimental design, and results.

6. **Chapter 6: Conclusion and Future Work**: This chapter summarizes the key findings of the thesis, highlights the main contributions, discusses limitations of the proposed approaches, and outlines potential future research directions, including extension to rupture risk prediction and multimodal clinical decision systems.

CHAPTER 2

Literature Review

2.1 Classical and Deep Learning-Based Detection in Medical Imaging

Early medical image segmentation methods were primarily based on intensity thresholding, region growing, active contour models, and deformable shape modeling. These methods relied heavily on handcrafted priors and were highly sensitive to noise and imaging artifacts, particularly in angiographic data where vessel overlap is common. With the advent of deep learning, CNN-based architectures significantly improved segmentation performance. The U-Net architecture became the de facto standard for biomedical segmentation due to its encoder-decoder structure and skip connections that preserve spatial details (15). Variants such as 3D U-Net extended these ideas to volumetric data, enabling improved performance on CTA and MRI datasets.

Subsequent works introduced attention mechanisms and transformer-based architectures to capture long-range dependencies. Models such as UNet (15), UNet++ (16), SwinUNetR (17), and TransUNet (18), which is particularly important in complex anatomical structures like cerebral vessels. However, despite these advancements, most segmentation approaches rely on pixel-wise losses such as Dice loss and cross-entropy that do not explicitly enforce boundary consistency, often resulting in coarse or blurred object edges. This limitation has motivated research into boundary-aware segmentation techniques, including level-set-inspired losses, edge supervision, and distance-transform-based methods, which are discussed further in Section 2.4.

2.2 Foundation Models in Medical Imaging

Foundation models have recently emerged as powerful general-purpose models trained on large-scale datasets, capable of transferring knowledge across diverse downstream tasks (8; 19; 20). In medical imaging, models such as SAM, MedSAM, and VISTA3D have demonstrated strong performance in segmentation and volumetric representation learning.

SAM introduced a promptable segmentation framework trained on large-scale natural image datasets, while MedSAM adapted this architecture to medical imaging domains. Similarly, 3D foundation models such as VISTA3D provide volumetric feature extraction capabilities for CT and MRI data, and have shown promising transferability to cerebrovascular imaging tasks when combined with domain-specific fine-tuning strategies.

Despite their success, foundation models (21; 22; 23) often lack domain-specific inductive biases required for precise medical interpretation. In particular, they do not inherently encode vascular topology or anatomical boundary constraints, which are crucial in neurovascular imaging tasks. This has led to increasing research interest in integrating structural and geometric priors into foundation model pipelines to improve reliability and clinical relevance, motivating both contributions of this thesis (24).

2.3 Aneurysm Detection and Vascular Structure Modeling

With the rise of deep learning, CNN-based detection models have been applied to CTA and MRA data for aneurysm classification and localization. Patch-based 3D CNNs (25) have demonstrated improved sensitivity but often suffer from high false positive rates due to lack of anatomical context. As illustrated in Fig. 2.1, conventional approaches formulate aneurysm detection as a simple binary classification task, identifying only the presence or absence of an aneurysm without any anatomical context regarding its location within the cerebrovascular tree (26; 27; 28; 29; 30).

Early investigations into automated aneurysm analysis relied on handcrafted geometric and intensity-based features extracted from vascular structures, coupled with classical classifiers. Among these, Malik et al. (12) proposes a 2D DSA-based framework for intracranial saccular aneurysm detection and rupture risk estimation using morphological analysis of cerebral angiograms. It segments vascular structures via watershed and distance transform, classifies aneurysms using an MLP with Haralick texture features, and performs geometric analysis to quantify rupture risk. Heit et al. (26) evaluated RAPID Aneurysm, a semi-automated AI system for aneurysm detection on CTA, reporting strong diagnostic performance. Dai et al. (27) introduced a deep learning pipeline based on Faster R-CNN applied to 2D projections derived from 3D CTA volumes, achieving high sensitivity particularly for aneurysms exceeding 3,mm in diameter while maintaining clinically viable inference times. Stember et al. (29) developed a U-Net-based convolutional architecture for automatic aneurysm detection and measurement from MRA images. Nakao et al. (30) presented a computer-aided detection system that combined CNNs with maximum intensity projection (MIP) representations for aneurysm identification in MRA data. Shi et al. (24) introduced a sham-AI experimental framework designed to evaluate the true clinical impact of AI-assisted aneurysm detection on CTA, finding that radiologists demonstrated robust-

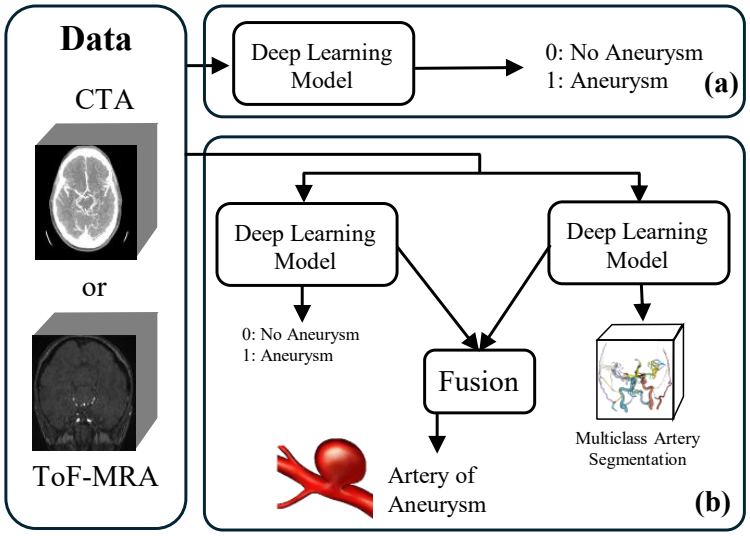

Figure 2.1: Comparison of (a) conventional binary aneurysm detection, which identifies only the presence or absence of an aneurysm, and (b) the proposed ARAN framework, which additionally identifies the artery of origin via a fused multiclass artery segmentation branch, enabling anatomically informed and clinically meaningful aneurysm localization.

ness to misleading AI suggestions and underscoring the necessity of developing reliable and interpretable AI tools for clinical integration. Cebeci et al. (31) advanced this direction by proposing a vessel-aware detection framework that employs multi-scale deformable 3D attention mechanisms for aneurysm localization in CT scans.

However, most existing approaches still treat aneurysm detection as a binary classification task without explicitly modeling the relationship between aneurysms and their parent arteries, as highlighted in Fig. 2.1(a). This limitation is clinically significant because treatment planning and rupture risk assessment depend heavily on the aneurysm's precise location within the vascular tree, including its parent artery, local branching geometry, and hemodynamic environment. The introduction of structure-aware modeling, as implemented in ARAN, addresses this gap by integrating vascular geometry descriptors and graph-based representations (13) of cerebral arteries with foundation model visual features, enabling the artery-level anatomical understanding shown in Fig. 2.1(b). Table 2.1 summarizes comparison of ARAN with the different proposed methods in literature, based on our experiments.

2.4 Boundary-Aware Learning in Medical Image Segmentation

In medical image segmentation, region-based losses such as Dice and Tversky (32) promote region-level overlap and mitigate class imbalance, while Focal loss (33) suppresses easy nega-

Table 2.1: Comparative analysis of existing intracranial aneurysm detection methods. ✓ = strength, ✗ = weakness.

Study	Modality	Artery-Aware	Architecture
Heit et al. (26)	CTA	✗	Semi-automated AI
Dai et al. (27)	CTA	✗	Faster R-CNN
Stember et al. (29)	MRA	✗	U-Net CNN
Nakao et al. (30)	MRA	✗	CNN + MIP
Shi et al. (24)	CTA	✗	CNN-based
Cebeci et al. (31)	CTA	✗	Deformable 3D Attn.
ARAN (Ours)	CTA + ToF-MRA	✓	VISTA3D + GNN

tives to redirect the model's attention toward harder, more informative examples. Despite their widespread adoption, these approaches operate purely on binary masks and remain insensitive to fine geometric details, frequently producing segmentations with incomplete or blurred boundaries. To address this, a number of distance and boundary-aware loss formulations have been proposed in the literature. Ma et al. (34) provided a comprehensive survey of segmentation methods, covering both architectural advances and loss function design. Karimi et al. (35) formulated a loss that directly minimizes the Hausdorff distance to penalize large contour deviations, while Kervadec et al. (36) proposed a boundary loss grounded in signed distance maps to improve contour localization, particularly for small anatomical structures. Xue et al. (37) took a different approach by directly regressing signed distance fields as the segmentation target. More recently, distance-based weighting schemes (38) have attracted growing interest as a means of incorporating geometry-aware supervision into deep learning pipelines. Despite these advances, conventional distance-based losses suffer from unstable gradient behavior by assigning uniform treatment to all pixels, leading to abrupt optimization updates and poor boundary convergence, especially in the early stages of training. In contrast, FocusSDF selectively emphasizes pixels in the vicinity of structural boundaries while exponentially suppressing contributions from distant regions, stabilizing the gradient signal for smoother convergence and sharper boundary delineation. Table 2.2 summarizes comparison of FocusSDF with the proposed losses in literature, based on our experiments.

2.5 Research Gaps

Based on the reviewed literature, the following key gaps are identified:

1. Absence of vascular structure-aware reasoning in aneurysm detection, with most methods treating detection as a binary classification task lacking anatomical context.

2. Lack of explicit boundary modeling in medical segmentation frameworks, particularly in the

Table 2.2: Comparison of distance transform and SDF-based loss functions with FocusSDF, based on our experiments. Abbreviations: H = High, M = Medium, L = Low.

Method	**Boundary Awareness**	**Representation Type**	**Segmentation Accuracy**	**Clinical Interpretation**	**Cross-Modality Generalization**
Xue et al. (37)	L	SDF	H	L	H
Karimi et al. (35)	M	DT	M	M	M
Kervadec et al. (36)	H	SDF	M	H	L
Shit et al. (38)	L	DT	M	L	H
FocusSDF (Ours)	H	SDF	H	H	H

context of high-resolution DSA aneurysm delineation.

3. Foundation models lacking domain-specific structural and anatomical constraints necessary for reliable neurovascular interpretation.

4. Weak cross-modal generalization across CTA, ToF-MRA, and DSA, with most methods designed for a single modality without consideration of the full clinical imaging pathway.

These gaps directly motivate the two contributions of this thesis, presented in the order of the clinical imaging pipeline:

1. **ARAN:** Introduces structure-aware aneurysm detection using vascular graph modeling and foundation model representations, enabling artery-level localization from non-invasive CTA and ToF-MRA scans.

2. **FocusSDF:** Addresses boundary-aware segmentation via signed distance function-based supervision, enabling precise aneurysm delineation from high-resolution DSA images to support morphological assessment and rupture risk quantification.

CHAPTER 3

Proposed Methodology

3.1 Overview of the Proposed Framework

This thesis presents a unified framework for intracranial aneurysm analysis that integrates structure-aware modeling and geometry-aware learning within a foundation model-driven paradigm, following the natural clinical imaging pathway. The proposed methodology is designed to address two complementary aspects of aneurysm analysis: (i) anatomically consistent localization of aneurysms within the cerebrovascular network from non-invasive CTA and ToF-MRA scans, and (ii) precise boundary delineation of the detected aneurysm sac from high-resolution DSA images to support morphological assessment and rupture risk quantification.

The overall framework consists of two primary components:

1. **ARAN: Structure-Aware Detection Module**
   This component models vascular topology and artery-specific geometry to enable anatomically informed aneurysm detection and localization using 3D foundation model representations from CTA and ToF-MRA scans. By explicitly identifying the parent artery of each detected aneurysm, ARAN provides clinically meaningful context for subsequent risk assessment.

2. **FocusSDF: Geometry-Aware Segmentation Module**
   This component focuses on learning fine-grained boundary representations of the detected aneurysm using signed distance function-based supervision on DSA images. It improves segmentation accuracy by emphasizing regions near anatomical boundaries, producing precise morphological delineations essential for rupture risk quantification.

While these modules operate on different modalities and dimensionalities, they share a common objective: enhancing foundation model representations through domain-specific structural and geometric priors to support the full aneurysm analysis pipeline.

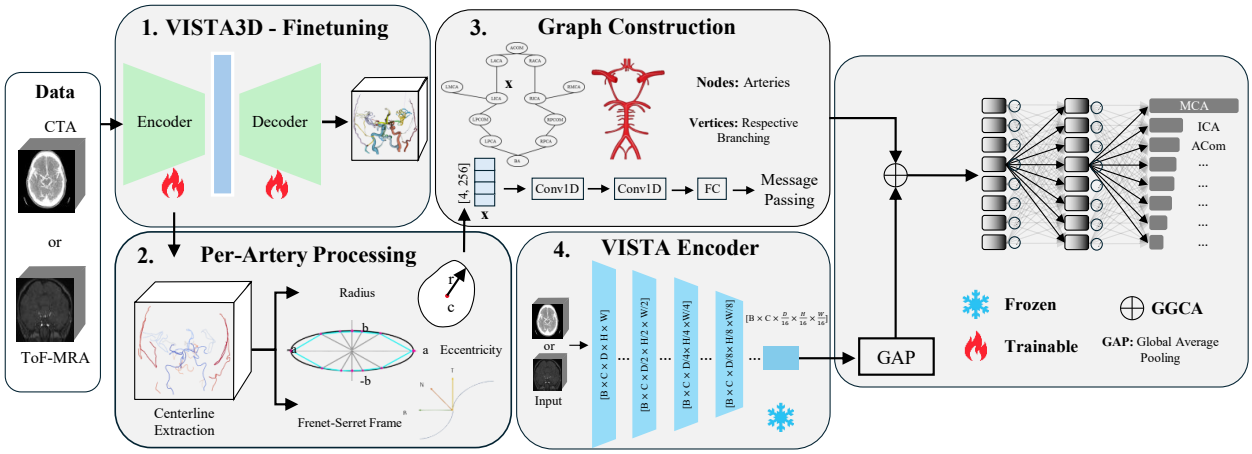

Figure 3.1: An overview of the proposed **ARAN** framework for artery-aware intracranial aneurysm detection. A 3D medical foundation model VISTA3D is first fine-tuned on CTA and ToF-MRA scans of normal subjects to perform multi-class artery segmentation, from which patient-specific geometric descriptors are derived. Volumetric image features are simultaneously extracted from the frozen VISTA3D encoder. The geometry-based graph representation and visual features are then fused through the proposed Geometry-Gated Cross-Attention (GGCA) module to identify the artery of origin of the detected aneurysm.

3.2 Structure-Aware Learning via Vascular Modeling

Our pipeline proceeds through multiple sequential stages, illustrated in Fig. 3.1 and elaborated in detail below.

3.2.1 Foundation Model Fine-Tuning for Multiclass Artery Segmentation

To obtain an anatomically consistent representation of the cerebrovascular network, we adopt VISTA3D (8), a medical vision foundation model for volumetric angiographic scans, to perform **multi-class artery segmentation** on CTA and ToF-MRA volumes of **normal subjects without aneurysms**. VISTA3D was selected over heavier transformer-based foundation models for three key reasons. First, its CNN-based SegResNet (39) encoder provides a favorable balance between computational efficiency and representational capacity, enabling volumetric feature extraction from 3D angiographic patches without prohibitive GPU memory requirements. Second, its support for both automatic and interactive segmentation modes makes it particularly well-suited for fine-tuning on multi-class artery segmentation under domain-specific supervision, where flexible adaptation to new anatomical targets is essential.

Let the input scan be denoted as:

$$\mathcal{X} \in \mathbb{R}^{H \times W \times D} \tag{3.1}$$

where H, W, D represent the spatial dimensions of the 3D scan.

The encoder f_θ produces a dense feature representation:

$$\mathcal{Z} = f_\theta(\mathcal{X}), \quad \mathcal{Z} \in \mathbb{R}^{C \times H' \times W' \times D'} \tag{3.2}$$

A segmentation head then predicts artery labels:

$$\hat{Y} = \text{Softmax}(g(\mathcal{Z})) \tag{3.3}$$

where $g(\cdot)$ denotes the decoder network and $\hat{Y}$ represents voxel-wise artery class probabilities. The model is optimized using cross-entropy and Dice losses to obtain accurate artery segmentation masks.

3.2.2 Vascular Geometry Feature Extraction

Once the arteries are extracted, we compute artery-specific centerlines using the open-source tool Kimimaro (40) and derive geometric descriptors that characterize the morphology of each vessel. These features encode structural information that may reflect abnormal vascular dilation or deformation associated with aneurysms. The computed geometric features are described below.

3.2.2.1 Radius

For each centerline point c_k, the vessel radius is estimated from the cross-sectional area A_k:

$$\varrho_k = \sqrt{\frac{A_k}{\pi}} \tag{3.4}$$

The resulting radius profile along the vessel is represented as:

$$\mathcal{R} = \{\varrho_1, \varrho_2, \ldots, \varrho_n\} \tag{3.5}$$

3.2.2.2 Cross-Sectional Eccentricity

To capture deviations from circular vessel shapes, eccentricity is computed using an ellipse fitted to each cross-section. Let α_k and ξ_k denote the major and minor axes, respectively:

$$\epsilon_k = \sqrt{1 - \frac{\xi_k^2}{\alpha_k^2}} \tag{3.6}$$

where $\epsilon_k \in [0, 1]$. Higher eccentricity indicates greater deviation from a circular cross-section, as illustrated in Fig. 1.1.

3.2.2.3 Frenet–Serret Curvature and Torsion

Let the vessel centerline be parameterized as a 3D curve:

$$\mathbf{c}(s) = (x(s), y(s), z(s)), \tag{3.7}$$

where s denotes arc length.

The Frenet–Serret frame defines an orthonormal basis along the curve consisting of the tangent ($\mathbf{T}$), normal ($\mathbf{N}$), and binormal ($\mathbf{B}$) vectors:

$$\mathbf{T}(s) = \frac{d\mathbf{c}}{ds}, \quad \mathbf{N}(s) = \frac{d\mathbf{T}/ds}{\left\|\frac{d\mathbf{T}}{ds}\right\|}, \quad \mathbf{B}(s) = \mathbf{T}(s) \times \mathbf{N}(s) \tag{3.8}$$

Curvature $\kappa(s)$ is defined as the rate of change of the tangent vector with respect to arc length:

$$\kappa(s) = \left\|\frac{d\mathbf{T}}{ds}\right\| \tag{3.9}$$

Torsion $\tau(s)$ measures the rate of change of the binormal vector and captures the twisting behavior of the curve:

$$\tau(s) = -\frac{d\mathbf{B}}{ds} \cdot \mathbf{N}(s) \tag{3.10}$$

These quantities are governed by the Frenet–Serret equations:

$$\frac{d}{ds}\begin{bmatrix}\mathbf{T}\\ \mathbf{N}\\ \mathbf{B}\end{bmatrix} = \begin{bmatrix}0 & \kappa & 0\\ -\kappa & 0 & \tau\\ 0 & -\tau & 0\end{bmatrix}\begin{bmatrix}\mathbf{T}\\ \mathbf{N}\\ \mathbf{B}\end{bmatrix} \tag{3.11}$$

These descriptors characterize vessel bending (via κ) and twisting (via τ), which serve as important geometric biomarkers potentially linked to aneurysm formation.

3.2.3 Vascular Graph Construction

To model inter-artery relationships and incorporate the discrepancies of arteries segmentation predicted by VISTA3D, the vascular structure is represented as a graph:

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}) \tag{3.12}$$

where nodes $\mathcal{V}$ correspond to individual arteries and edges $\mathcal{E}$ encode anatomical connections or spatial proximity.

Each node u_k is associated with aggregated geometric features:

$$\mathbf{h}_k = [\bar{\varrho}_k,\ \bar{\epsilon}_k,\ \bar{\kappa}_k,\ \bar{\tau}_k] \tag{3.13}$$

A Graph Attention Network propagates information across the vascular network as follows:

$$\mathbf{h}_k^{(l+1)} = \sigma\left(\sum_{m \in \mathcal{N}(k)} \mathbf{W}^{(l)} \mathbf{h}_m^{(l)}\right) \tag{3.14}$$

where $\mathcal{N}(k)$ denotes the set of neighboring arteries of node k and σ is a nonlinear activation function.

The choice of a Graph Neural Network for vascular structure modeling is motivated by the inherently relational nature of the cerebrovascular system. The Circle of Willis forms a topological graph in which arterial segments are functionally and anatomically connected, and meaningful clinical reasoning about aneurysm location requires modeling these inter-arterial relationships explicitly. A GNN achieves this through iterative message passing, allowing each artery node to aggregate geometric context from its neighbors across the vascular graph.

Furthermore, intracranial aneurysms frequently occur at vascular bifurcations and can be visually similar to arterial loops or regions of high vessel tortuosity in angiographic images. By propagating geometry-rich descriptors - including radius profiles, eccentricity, and Frenet-Serret curvature and torsion - across the graph, the GNN learns neighborhood-aware representations that aid in distinguishing pathological vascular dilations from normal anatomical complexity. This capability is not achievable by per-artery feedforward networks that process each vessel in isolation.

Additionally, since the geometry branch operates on centerline-derived descriptors rather than raw voxel-level segmentation outputs, it remains robust to local imperfections in the VISTA3D artery segmentation masks, providing a topologically grounded and segmentation-error-resilient complementary signal to the visual feature stream.

3.2.4 Visual Feature Extraction and Geometry-Gated Cross-Modal Fusion

In parallel with vascular graph modeling, the frozen VISTA3D encoder extracts deep volumetric features of the angiographic scan from the bottleneck. Let:

$$\mathcal{X} \in \mathbb{R}^{H \times W \times D}, \qquad \mathbf{f}_{\text{vis}} = f_\theta(\mathcal{X}) \in \mathbb{R}^{C \times H' \times W' \times D'} \tag{3.15}$$

We define an artery-specific region of interest (ROI) $\mathcal{R} \subset \Lambda$, where Λ denotes the voxel grid.

A pooled feature descriptor is obtained via ROI-restricted global average pooling:

$$\mathbf{z}_{\text{vis}} = \frac{1}{|\mathcal{R}|} \sum_{\mathbf{p} \in \mathcal{R}} \mathbf{f}_{\text{vis}}(\mathbf{p}) \ \in \ \mathbb{R}^C. \tag{3.16}$$

A learnable projection maps this to a compact visual embedding:

$$\mathbf{e}_{\text{vis}} = \mathbf{W}_v \mathbf{z}_{\text{vis}} \ \in \ \mathbb{R}^d. \tag{3.17}$$

Geometry-Gated Cross-Attention (GGCA) Let $\mathbf{e}_{\text{geo}} \in \mathbb{R}^d$ denote the geometry-aware d-dimensional embedding derived from the vascular graph encoder.

We define the following linear mappings:

$$\mathbf{Q} = \mathbf{W}_Q \, \mathbf{e}_{\text{geo}}, \quad \mathbf{K} = \mathbf{W}_K \, \mathbf{e}_{\text{vis}}, \quad \mathbf{V} = \mathbf{W}_V \, \mathbf{e}_{\text{vis}}, \tag{3.18}$$

with $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^d$.

The cross-attended representation is computed as:

$$\mathbf{a} = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}} \right) \mathbf{V} \ \in \ \mathbb{R}^d. \tag{3.19}$$

To modulate the contribution of visual features based on geometric reliability, we introduce a geometry-conditioned gating function:

$$\mathbf{g} = \sigma(\mathbf{W}_g \, \mathbf{e}_{\text{geo}}) \ \in \ (0, 1)^d. \tag{3.20}$$

The fused embedding is then defined as:

$$\mathbf{e}_{\text{out}} = (1 - \mathbf{g}) \odot \mathbf{e}_{\text{geo}} + \mathbf{g} \odot \mathbf{a}, \tag{3.21}$$

where $\odot$ denotes element-wise multiplication, as illustrated in Fig. 3.2.

Finally, the representation is refined via a residual feed-forward transformation:

$$\mathbf{e}_{\text{out}} = \text{LayerNorm}(\mathbf{e}_{\text{out}} + \phi_f(\mathbf{e}_{\text{out}})) \, , \tag{3.22}$$

where ϕ_f denotes a two-layer position-wise feed-forward network.

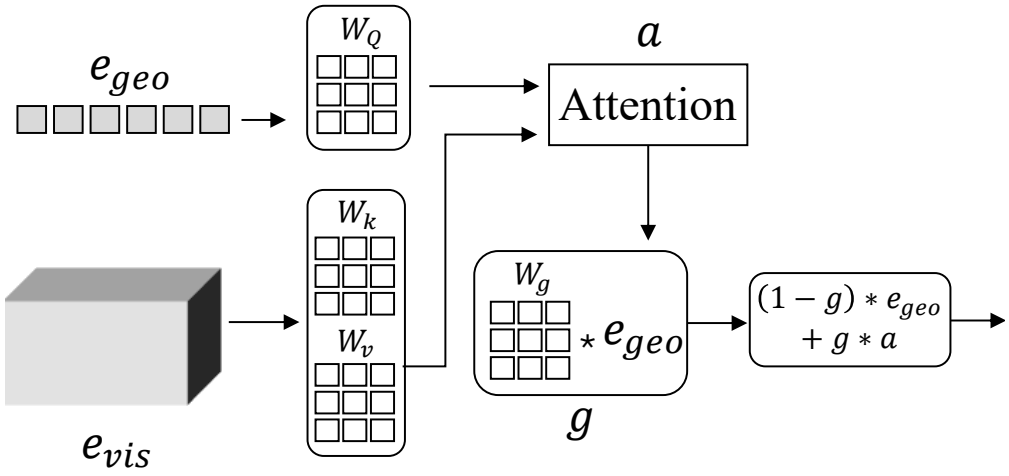

Figure 3.2: Detailed illustration of the proposed GGCA module. The geometry embedding $\mathbf{e}_{\text{geo}}$ derived from the vascular graph encoder serves as the query, while the visual embedding $\mathbf{e}_{\text{vis}}$ from the VISTA3D encoder provides the keys and values. A geometry-conditioned gate $\mathbf{g}$ adaptively controls the contribution of visual features relative to structural priors, enabling context-aware fusion of appearance and vascular geometry for artery-of-origin prediction.

3.3 Geometry-Aware Learning via Signed Distance Functions

Once an aneurysm is detected and its parent artery identified via ARAN, the patient is referred for a high-resolution DSA scan to enable detailed morphological assessment. FocusSDF is then applied to precisely delineate the aneurysm boundary from the DSA image, providing the fine-grained segmentation required for rupture risk quantification.

3.3.1 Signed Distance Function

Let $\mathcal{D} \subset \mathbb{R}^n$ denote a region with boundary $\partial\mathcal{D}$. For any point $\mathbf{q} \in \mathbb{R}^n$, the Euclidean distance from $\mathbf{q}$ to the nearest point on the boundary $\partial\mathcal{D}$ is formally defined as:

$$\rho(\mathbf{q}, \partial\mathcal{D}) = \min_{\mathbf{b}\in\partial\mathcal{D}} \|\mathbf{q} - \mathbf{b}\|_2 \tag{3.23}$$

where $\|\cdot\|_2$ denotes the Euclidean norm and the minimization is taken over all boundary points $\mathbf{b} \in \partial\mathcal{D}$. The *SDF* is then defined as a signed extension of this distance, encoding both the proximity to the boundary and the spatial relationship of each point with respect to the region interior:

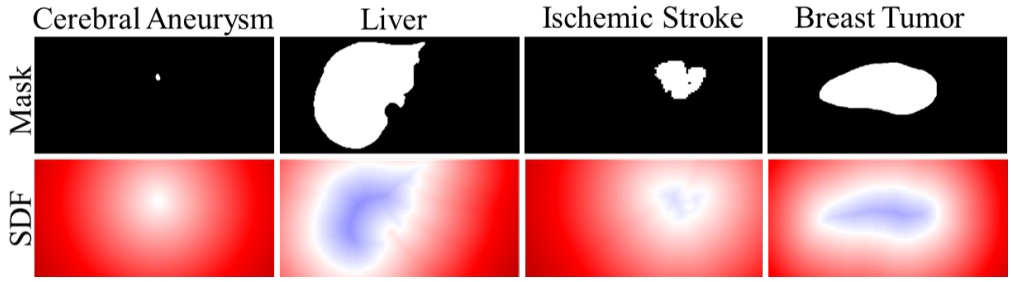

Figure 3.3: SDF representations of binary segmentation masks across different anatomical regions. Negative values indicate interior pixels, zero values correspond to the object boundary, and positive values represent exterior pixels. The magnitude at each location encodes the Euclidean distance to the nearest boundary point, providing a continuous and geometry-rich supervision signal for boundary-aware learning.

$$\psi(\mathbf{q}) = \begin{cases} -\rho(\mathbf{q}, \partial\mathcal{D}), & \text{if } \mathbf{q} \in \mathcal{D}, \\ 0, & \text{if } \mathbf{q} \in \partial\mathcal{D}, \\ \rho(\mathbf{q}, \partial\mathcal{D}), & \text{if } \mathbf{q} \in \mathcal{D}^c \end{cases} \tag{3.24}$$

By convention, interior points $\mathbf{q} \in \mathcal{D}$ receive negative values, boundary points $\mathbf{q} \in \partial\mathcal{D}$ are assigned a value of zero, and exterior points $\mathbf{q} \in \mathcal{D}^c$ carry positive values. The magnitude $|\psi(\mathbf{q})|$ at any location reflects the shortest distance to the boundary, providing a continuous and geometrically meaningful representation of the region's shape. Fig. 3.3 illustrates the SDF representation of binary masks across different anatomical regions, demonstrating how the signed distance encodes both interior structure and boundary proximity in a unified and spatially coherent manner.

3.3.2 Proposed Loss Function Formulation

To prioritize learning in the vicinity of object boundaries, we introduce an adaptive, distance-based weighting mechanism for SDF supervision. Unlike binary mask losses that assign uniform importance to all pixels, FocusSDF allocates higher weights to pixels near the boundary, where geometric precision is most critical.

Let $\mathcal{F}$ and $\hat{\mathcal{F}}$ denote the ground truth (GT) and predicted SDF maps, respectively, and let $\mathcal{F}_j$ and $\hat{\mathcal{F}}_j$ represent the GT and predicted SDF values at the j-th pixel.

The spatial domain $\mathcal{D}$ is partitioned as:

$$\mathcal{D}_{\text{in}} = \{\, \mathbf{j} : \mathcal{F}_{\mathbf{j}} < 0 \,\}, \quad \partial\mathcal{D} = \{\, \mathbf{j} : \mathcal{F}_{\mathbf{j}} = 0 \,\}, \quad \mathcal{D}_{\text{out}} = \{\, \mathbf{j} : \mathcal{F}_{\mathbf{j}} > 0 \,\}$$

The proposed FocusSDF loss integrates the boundary-aware weighting with a gradient consis-

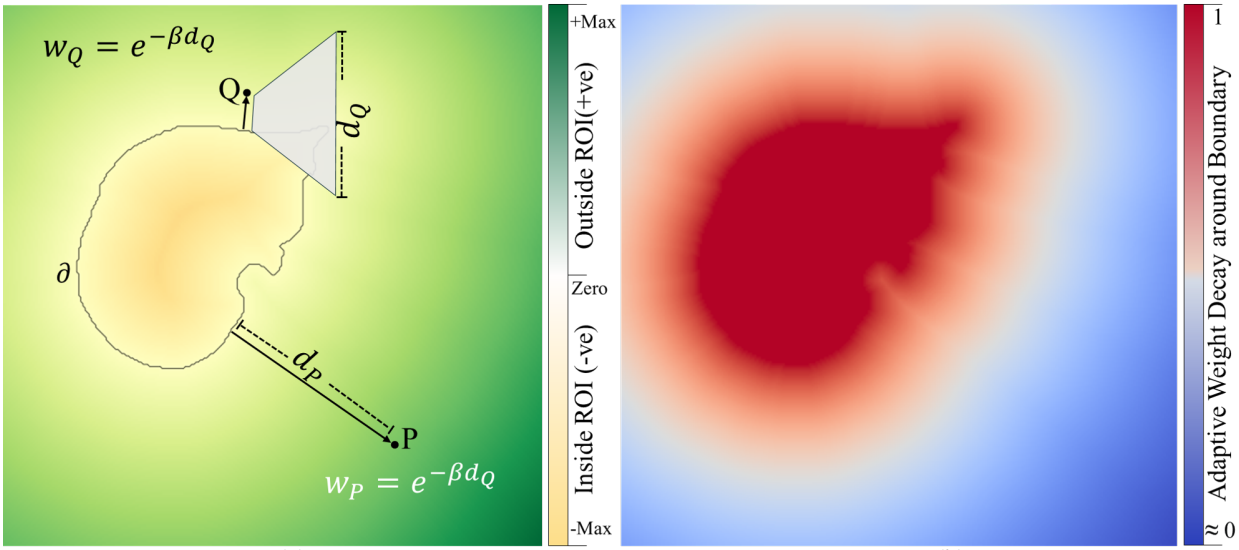

(a) (b)

Figure 3.4: Visualization of the proposed boundary-aware weighting mechanism in **FocusSDF**. (a) Signed distance map showing two representative pixels: point P, located far from the boundary, and point Q, located in close proximity to it. Darker green or yellow shading indicates greater distance from the ROI boundary. (b) Corresponding spatial weight heatmap illustrating elevated loss contributions near the boundary and exponentially diminishing weights for distant exterior pixels, shown here for $\beta = 0.005$.

tency term:

$$\mathcal{L}_{\text{FocusSDF}} = \frac{1}{|\mathcal{D}|} \sum_{j \in \mathcal{D}} v_j \cdot \left| \mathcal{F}_j - \hat{\mathcal{F}}_j \right|^p + \mu \frac{1}{|\mathcal{D}|} \sum_{j \in \mathcal{D}} \left| \nabla \mathcal{F}_j - \nabla \hat{\mathcal{F}}_j \right|^p, \tag{3.25}$$

where,

$$v_j = \begin{cases} 1, & \text{if } \mathbf{j} \in \mathcal{D}_{\text{in}} \text{ or } \partial\mathcal{D} \\ \exp\left(-\beta \cdot \mathcal{F}_{\mathbf{j}}\right), & \text{if } \mathbf{j} \in \mathcal{D}_{\text{out}} \end{cases}$$

where the gradient consistency term computes the difference between gradient map norms ∇, $p \in \{1, 2\}$, $\beta > 0$ governs the rate of weight decay away from the object boundary, and μ balances the two terms. Fig. 3.4 demonstrates the behavior of the proposed weight decay mechanism in signed distance maps for points P and Q, located far from and relatively close to the boundary of the region of interest (ROI), respectively. In Fig. 3.4(a), darker shades of green or yellow indicate greater distance from the ROI boundary. In Fig. 3.4(b), darker red shading reflects higher pixel weights in terms of loss contribution.

where:

- $p \in \{1, 2\}$
- β governs spatial decay

- μ balances gradient consistency

3.4 Unified Insight

The proposed framework establishes a clinically grounded and structurally coherent pipeline for intracranial aneurysm analysis by integrating anatomically informed vascular reasoning with precise lesion boundary delineation. Following the natural clinical imaging pathway, ARAN first localizes the aneurysm and identifies its parent artery from non-invasive CTA and ToF-MRA scans, providing the structural context essential for risk stratification and treatment planning. Subsequently, FocusSDF enables fine-grained aneurysm boundary delineation from high-resolution DSA images, producing morphologically precise segmentations that directly support quantitative rupture risk assessment. By jointly leveraging vascular structure and lesion geometry through complementary foundation model representations, the framework moves beyond isolated aneurysm detection toward a unified, clinically meaningful pipeline that captures both global anatomical relationships and local morphological characteristics, laying the groundwork for fully automated artery-aware rupture risk prediction.

CHAPTER 4

Datasets

This chapter describes the datasets utilized across the two components of the proposed framework. Three categories of datasets are employed: (i) a dataset for foundation model fine-tuning for multi-class artery segmentation, (ii) datasets for artery-aware aneurysm detection, and (iii) datasets for boundary-aware aneurysm segmentation. Together, these datasets span four imaging modalities - CTA, ToF-MRA, DSA, MRI, and Ultrasound - and cover a diverse range of anatomical structures and clinical scenarios, enabling comprehensive evaluation of both ARAN and FocusSDF across realistic and challenging settings.

4.1 Datasets for ARAN: Artery-Aware Aneurysm Detection

4.1.1 TopCoW 2024: Foundation Model Fine-Tuning

For multi-class artery segmentation, we fine-tune the VISTA3D foundation model using the **TopCoW 2024** dataset (41). This dataset was released as part of the Topology-Aware Segmentation of the Circle of Willis challenge and represents one of the few publicly available resources providing paired CTA and ToF-MRA volumes with detailed cerebrovascular annotations. It comprises 125 scans across both modalities, making it particularly well-suited for multi-modal foundation model adaptation in the neurovascular domain.

The dataset provides annotations for the major arterial segments of the Circle of Willis, from which we select six arterial locations most commonly associated with intracranial aneurysm occurrence: the ICA, ACom, MCA, PCom, BA, and PCA. Given that the primary objective of this study is to develop an artery-aware aneurysm rupture risk prediction framework, and based on clinical insights indicating comparable rupture risk profiles between left and right counterparts of the same arterial structure, we merge bilateral artery annotations into unified classes. This consolidation results in six anatomically consistent and clinically meaningful artery categories that serve as the target classes for downstream analysis and artery-of-origin identification.

Since explicit artery-of-origin labels for aneurysms are not directly available in this dataset,

aneurysm candidates are first localized using zero-shot predictions from the VISTA3D model. The corresponding artery-of-origin labels are subsequently extracted from the segmented arterial anatomy and verified by collaborating neurosurgeons at Henry Ford Health System to ensure clinical correctness, using 3D reconstructed arterial models as reference. This annotation pipeline ensures that the labels used for downstream training are both anatomically accurate and clinically validated.

4.1.2 OpenNeuro 2022 and OpenNeuro 2024: IA Detection in ToF-MRA

For artery-aware aneurysm origin prediction on ToF-MRA data, we combine two publicly available datasets: **OpenNeuro 2022** (42) and **OpenNeuro 2024** (43). Both datasets provide Time-of-Flight Magnetic Resonance Angiography volumes with aneurysm annotations, offering a non-invasive imaging perspective on cerebrovascular pathology. ToF-MRA is particularly advantageous for aneurysm screening due to its high vessel-to-background contrast without the need for contrast agent administration, though it introduces challenges such as flow artifacts and lower spatial resolution compared to CTA.

To maintain consistency in the experimental protocol and avoid confounding factors associated with multiple aneurysms per patient, we restrict our analysis to cases containing a single aneurysm per subject. Following this filtering criterion, the combined dataset yields a total of **157 subjects**, distributed across three artery classes: 59 cases with aneurysms originating from the Internal Carotid Artery (ICA), 47 from the Anterior Communicating Artery (ACom), and 51 from the Middle Cerebral Artery (MCA). These three arterial locations represent the most clinically prevalent sites of intracranial aneurysm occurrence and collectively account for the vast majority of cases encountered in neurovascular practice.

4.1.3 Large Intracranial Aneurysm Segmentation Dataset: IA Detection in CTA

For artery-aware aneurysm origin prediction on CTA data, we employ the **Large Intracranial Aneurysm Segmentation Dataset** (44). This dataset provides three-dimensional CTA volumes with detailed aneurysm segmentation annotations, making it one of the larger publicly available resources for intracranial aneurysm analysis in the CTA modality. CTA offers high spatial resolution and is widely used in clinical practice for aneurysm detection and follow-up, though it involves ionizing radiation and contrast agent administration.

Following the same single-aneurysm-per-patient protocol applied to the ToF-MRA datasets, we construct a balanced subset of **200 subjects**, carefully selected to ensure approximately equal representation across the three target artery classes: 66 cases with ICA aneurysms, 67 with ACom

aneurysms, and 66 with MCA aneurysms. This class-balanced construction is intentional, as it mitigates the influence of class frequency imbalance on classification performance and ensures that evaluation metrics reflect genuine discriminative ability across all artery classes rather than majority-class bias.

4.2 Datasets for FocusSDF: Boundary-Aware Segmentation

To evaluate the generalizability of the proposed FocusSDF loss function across diverse anatomical structures and imaging modalities, experiments are conducted on four datasets spanning DSA, CTA, MRI, and Ultrasound. This multi-dataset, multi-modality evaluation is intentional: it demonstrates that the boundary-aware weighting formulation of FocusSDF is architecture-agnostic and modality-independent, rather than being tailored to a single imaging domain.

For each dataset, a consistent data partitioning protocol is applied: 20% of patients are allocated to the held-out test set, while 5% of the original training split is further reserved for validation. This ensures that test set evaluation is performed on entirely unseen data, providing an unbiased estimate of generalization performance.

4.2.1 In-House Cerebral Aneurysm Dataset: DSA

The primary dataset for aneurysm segmentation evaluation is an in-house collection of **50 DSA patient studies**, acquired at Henry Ford Health System under Institutional Review Board approval (IRB No. 11254). Digital Subtraction Angiography represents the gold standard modality for intracranial vascular assessment, offering the highest spatial resolution among available angiographic techniques. However, DSA images present significant segmentation challenges due to the complex overlapping of arterial loops, aneurysm sacs, and surrounding vascular structures, all of which share similar dye contrast intensities, as illustrated in Fig. 1.2(a).

The test subset was intentionally composed of the most clinically complex cases in the collection, including challenging projection angles, partially occluded aneurysms, and multi-aneurysm patients. This deliberate selection was made to stress-test the boundary delineation capability of competing methods under realistic clinical conditions, rather than evaluating on straightforward cases that may not reflect deployment scenarios. Binary aneurysm segmentation masks were prepared and clinically verified in collaboration with Henry Ford Health System.

4.2.2 Liver Segmentation Dataset: CTA

To assess FocusSDF on a well-established large-organ segmentation benchmark, we employ the publicly available **liver segmentation dataset** (45) acquired in the CTA modality. Liver segmen-

tation in CTA presents challenges related to low tissue contrast at organ boundaries and partial overlap with adjacent soft tissue structures, as illustrated in Fig. 1.2(b). This dataset serves as a representative benchmark for evaluating boundary precision in a high-contrast volumetric setting, complementing the fine-grained boundary challenges posed by the DSA aneurysm dataset.

4.2.3 Ischemic Stroke Segmentation Dataset: MRI

For evaluation on MRI data, we utilize the publicly available **ISLES (Ischemic Stroke Lesion Segmentation)** dataset (46). Ischemic stroke lesion segmentation in MRI is particularly challenging due to the highly irregular and diffuse lesion boundaries, combined with near-identical intensity profiles between infarcted and surrounding healthy tissue, as illustrated in Fig. 1.2(c). This dataset tests the ability of FocusSDF to enforce boundary sensitivity in a low-contrast, morphologically complex setting that is fundamentally different from the angiographic and CT domains.

4.2.4 Breast Tumor Segmentation Dataset: Ultrasound

Finally, to evaluate FocusSDF on a non-radiological modality, we employ the publicly available **breast ultrasound tumor segmentation dataset** (47). Ultrasound imaging introduces unique challenges including speckle noise, acoustic shadowing, and significant structural similarity between tumor tissue and surrounding healthy parenchyma, as illustrated in Fig. 1.2(d). This dataset extends the evaluation of FocusSDF to a clinically important modality that is widely used for breast cancer screening, demonstrating the cross-modal generalization capacity of the proposed boundary-aware supervision strategy.

4.3 Dataset Summary

Table 4.1 provides a consolidated overview of all datasets used in this study, including modality, dataset size, clinical task, and usage context within the proposed framework.

Table 4.1: Summary of datasets used in this study across both framework components.

Dataset	**Modality**	**# of Subjects**	**# of Slices**	**Task**	**Component**
TopCoW 2024 (41)	CTA + ToF-MRA	125	–	Artery Segmentation	ARAN
OpenNeuro 2022 (42)	ToF-MRA	115	–	Aneurysm Detection	ARAN
OpenNeuro 2024 (43)	ToF-MRA	42	–	Aneurysm Detection	ARAN
Large IA Dataset (44)	CTA	200	–	Aneurysm Detection	ARAN
In-House IA Dataset	DSA	50	$\sim$3800	Aneurysm Segmentation	FocusSDF
Liver Dataset (45)	CTA	20	$\sim$2400	Liver Segmentation	FocusSDF
ISLES (46)	MRI	250	$\sim$900	Stroke Segmentation	FocusSDF
Breast Ultrasound (47)	Ultrasound	650	$\sim$650	Tumor Segmentation	FocusSDF

CHAPTER 5

Experimental Setup and Results

This chapter presents the experimental evaluation of both proposed components of the unified framework. Following the clinical motivation outlined throughout this thesis, we first evaluate ARAN on aneurysm detection and artery-of-origin identification from non-invasive CTA and ToF-MRA scans, and subsequently evaluate FocusSDF on precise aneurysm boundary delineation from DSA scans alongside other medical imaging modalities.

5.1 ARAN: ARtery-Aware Intracranial ANeurysm Detection

5.1.1 Implementation Details

All experiments were implemented in Python 3.10.18 using PyTorch 2.10.0+cu128 (48). Training and inference were performed on a single NVIDIA GPU with 48 GB memory on the University of Michigan's Great Lakes HPC cluster, equipped with 372 GB system RAM. Input patches of $96\times96\times96$ voxels were extracted prior to training.

For foundation model fine-tuning, the VISTA3D encoder was initialised with publicly available pretrained weights and optimised for six-class artery segmentation on CTA and ToF-MRA volumes using DiceCE loss, AdamW optimiser with learning rate 10^{-5}, and early stopping over 175 epochs.

For ARAN training, the AdamW optimiser was used with learning rate 3×10^{-4}, weight decay 10^{-3}, and cosine annealing decay to 10^{-6}, with gradient clipping at 1.0. The model was trained for 296 epochs with early stopping and an 80/20 stratified split. The GAT comprised 3 layers with 4 attention heads and hidden dimension 64 over a 6-node Circle-of-Willis graph. Vision features were projected from the 768-dimensional VISTA3D bottleneck to 256 dimensions via a linear layer with LayerNorm and ReLU. Class-weighted cross-entropy loss was used to handle class imbalance. For vision-only baselines, global average pooling was applied over the encoder features and passed through two fully connected layers of 256 neurons each with ReLU activations, followed by a linear projection to 3 output classes.

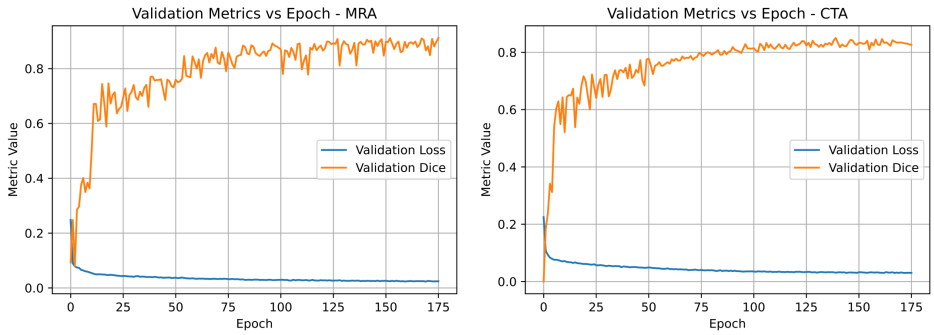

Figure 5.1: Fine-tuning dynamics of VISTA3D for multi-class artery segmentation on CTA and ToF-MRA modalities. The plots show the evolution of validation loss and per-class Dice score across training epochs, demonstrating stable convergence and consistent improvement for all six artery classes in both modalities.

5.1.2 Evaluation Metrics

We assess both arterial segmentation quality and artery-aware aneurysm detection performance. For multiclass artery segmentation, we report the Dice Similarity Coefficient (DSC) per artery class and its macro-average across all six classes:

$$\text{Dice} = \frac{2|P \cap G|}{|P| + |G|} \tag{5.1}$$

where P and G denote the predicted and ground truth binary masks, respectively.

For artery-aware aneurysm detection, we report overall accuracy and macro F1-score:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{5.2}$$

$$\text{F1}_{\text{macro}} = \frac{1}{C} \sum_{c=1}^{C} \text{F1}_c \tag{5.3}$$

where C denotes the number of artery classes. Macro F1 assigns equal weight to each class regardless of class frequency.

5.1.3 Results

5.1.3.1 Zero-Shot Performance of VISTA3D

We first evaluated VISTA3D's ability to segment cerebral arteries in a zero-shot setting on the TopCoW dataset. Across multiple artery classes, the model exhibited near-zero Dice scores, indicating that its pretrained knowledge did not generalize to this domain. We identify three contributing fac-

tors:

1. **Domain Mismatch**: As VISTA3D was pretrained exclusively on abdominal organs with no exposure to cerebrovascular structures.

2. **Structural Complexity**: As arteries form hierarchical branching trees with complex topological dependencies that differ fundamentally from isolated organs.

3. **Resolution Mismatch**: As the pretrained model was trained on images with voxel spacing of $(1.5, 1.5, 1.5)$ mm, whereas cerebral artery datasets have finer resolution with an average spacing of $(0.4, 0.4, 0.7)$ mm.

5.1.3.2 Fine-Tuned VISTA3D for Artery Segmentation

Table 5.1 reports Dice scores for both CTA and ToF-MRA modalities after fine-tuning. Despite VISTA3D being predominantly pretrained on CTA data, fine-tuning on ToF-MRA yielded substantially higher Dice scores. This is attributed to the significant intensity contrast between arteries and background in ToF-MRA compared to CTA. The fine-tuning dynamics, including validation loss and Dice score trends for both modalities, are illustrated in Fig. 5.1, demonstrating stable convergence across epochs.

Table 5.1: Validation Results of VISTA3D fine-tuning for six-class artery segmentation on TopCow 2024 Dataset.

	ToF-MRA		CTA	
Artery Class	Dice	IoU	Dice	IoU
ICA	0.9490	0.9029	0.8529	0.7433
ACom	0.8596	0.7536	0.7219	0.5650
MCA	0.9302	0.8697	0.8073	0.6772
PCom	0.8744	0.7763	0.7513	0.6016
BA	0.9320	0.8726	0.8328	0.7136
PCA	0.9070	0.8306	0.7827	0.6429
Mean	0.9287	0.8671	0.8415	0.7260

5.1.3.3 Artery-Aware Aneurysm Detection

Table 5.2 presents the quantitative comparison across vision-only, geometry-only, and hybrid methods. The vision-only baseline feeds VISTA3D encoder features into a simple classifier to predict the artery of origin. The geometry-only baseline relies exclusively on vascular geometric descriptors derived from vessel centerlines. The proposed ARAN framework integrates visual represen-

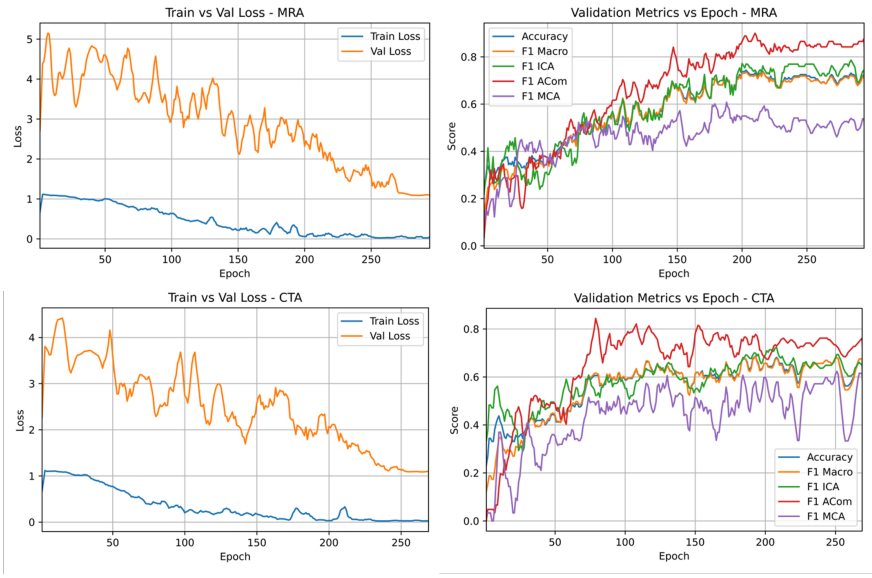

Figure 5.2: Training and validation performance of ARAN for artery-aware aneurysm detection on CTA and ToF-MRA datasets. Left: training and validation loss curves for both modalities. Right: validation metrics across epochs, including overall accuracy, macro F1-score, and class-wise F1-scores for ICA, ACom, and MCA. Results demonstrate steady performance improvement and stable convergence with consistent gains across all artery classes.

tations from the VISTA3D encoder with physics-inspired vascular graph features, leveraging both appearance and artery-specific geometric context.

The results demonstrate that integrating vascular geometry with visual representations leads to improved artery-level classification performance compared to either modality alone. Training and validation behavior, shown in Fig. 5.2, demonstrate steady improvements and stable convergence across both CTA and ToF-MRA datasets. t-SNE visualizations of the learned embeddings, shown in Fig. 5.3, further confirm that ARAN learns discriminative representations with well-separated artery-class clusters, validating the effectiveness of the geometry-gated cross-attention fusion strategy.

5.2 FocusSDF: Geometry-Aware Aneurysm Segmentation

A critical yet often overlooked factor in achieving this precision is the choice of loss function: performance is driven as much by the loss as the model - poor losses limit strong models, while effective losses boost simpler ones. Following aneurysm detection and artery-of-origin identification via ARAN, the detected aneurysm region serves as the basis for precise morphological delineation. Once a clinician orders a DSA scan following ARAN-based detection, FocusSDF is applied to segment the aneurysm sac with high boundary precision, providing the anatomically

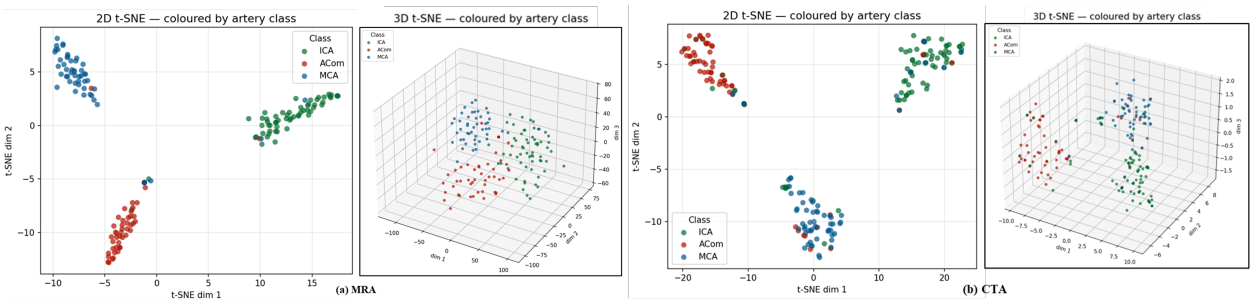

Figure 5.3: t-SNE visualization of learned feature embeddings from ARAN for artery-of-origin classification. Both 2D and 3D projections are shown for CTA and ToF-MRA datasets. Distinct and well-separated clustering of ICA, ACom, and MCA classes demonstrates that ARAN learns highly discriminative representations, confirming the effectiveness of geometry-gated cross-attention fusion in isolating artery-specific identity signals.

Table 5.2: Comparison of vision-only, geometry-only, and hybrid methods on ToF-MRA and CTA datasets. Acc = overall accuracy, $F1_{mac}$ = macro F1, Bold = best in column, GAP = Global Average Pooling, MLP = Multi-Layer Perceptron, † = Our method.

Modality	**Method**	**ToF-MRA**		**CTA**	
		Acc	$F1_{mac}$	Acc	$F1_{mac}$
Vision	ResNet-50 (49)	0.581	0.563	0.537	0.528
	DenseNet-121 (50)	0.623	0.619	0.594	0.571
	VISTA3D	0.592	0.588	0.541	0.525
Geometry	GNN + MLP †	0.681	0.665	0.624	0.618
Hybrid	Ceballos et al. (31)	0.724	–	0.698	–
	Proposed: ARAN†	**0.858**	**0.836**	**0.782**	**0.753**

consistent delineations required for quantitative morphological assessment and artery-aware rupture risk prediction.

5.2.1 Datasets and Compared Loss Functions

Experiments were conducted on an in-house cerebral aneurysm dataset of 50 patients acquired in DSA modality, and three publicly available datasets: liver segmentation (CTA) (45), ischemic stroke segmentation (MRI) (46), and breast tumor segmentation (Ultrasound) (47). For each dataset, 20% of patients were allocated to the test set, while 5% of the original training split was reserved for validation. The in-house aneurysm test set, collected from Henry Ford Health System (IRB No. 11254), was intentionally composed of highly complex cases including challenging projection views and multi-aneurysm patients, to reflect real-world clinical settings and stress-test boundary delineation performance.

FocusSDF is benchmarked against Dice loss and four SDF-based losses from the literature (37; 35; 36; 38). Each referenced loss was combined with Dice loss (34) to stabilize training, and the resulting formulations are denoted $\mathcal{L}_1$, $\mathcal{L}_2$, $\mathcal{L}_3$, and $\mathcal{L}_4$, respectively.

5.2.2 Implementation Details

All experiments were implemented in Python 3.13.1 and PyTorch 2.6.0 with CUDA 11.8 on an NVIDIA RTX 4090 GPU. Input images were resized to 256×256 and intensity-normalized prior to training. Ground-truth signed distance fields were computed from binary masks and independently z-score normalized per image to zero mean and unit variance. All networks predict SDFs using a regression head with linear output (no activation), and during inference, predicted SDFs are thresholded at the zero level set to obtain binary segmentation masks. Training was conducted for 200 epochs with a batch size of 8 using AdamW (lr $= 10^{-4}$). The decay parameter β was selected via ablation over the range $[0.001, 0.01]$, with $\beta = 0.005$ yielding the best performance on the aneurysm segmentation dataset. We set $p = 2$ for smoothness and differentiability, and $\mu = 1$ to assign equal importance to the SDF regression and gradient consistency terms.

Backbones evaluated include UNet, UNet++, SwinUNetR, and TransUNet, covering both convolutional and transformer-based paradigms. Performance is assessed using three complementary metrics: **Dice** and **IoU** for region overlap quality, and **HD95** (95th percentile Hausdorff distance, in pixels) for boundary accuracy.

5.2.3 Results

Table 5.3 summarizes the quantitative results across all datasets and backbone architectures. FocusSDF achieves consistently superior overlap and boundary accuracy, most notably on the cerebral aneurysm dataset where boundary precision is most clinically critical. Its performance gains across both CNN and transformer-based architectures highlight its generality and robustness across diverse anatomical structures and imaging modalities. Fig. 5.4 shows that FocusSDF yields smoother validation curves and reduced performance variance across epochs compared to Dice loss and competing SDF-based losses, indicating more stable optimization and improved convergence behavior. Qualitative segmentation comparisons across all four datasets are presented in Fig. 5.5.

To complement these results, we additionally fine-tuned the foundation model **MedSAM** on the aneurysm dataset using Dice loss. Despite fine-tuning, it achieved a modest Dice score of 0.3340 on DSA images, likely due to the highly intricate vascular structure and overlapping dye intensities of aneurysms, arterial loops, and the surrounding vascular network, compounded by the complete absence of DSA modality in its pretraining data. Nevertheless, its one-shot generalization on CT

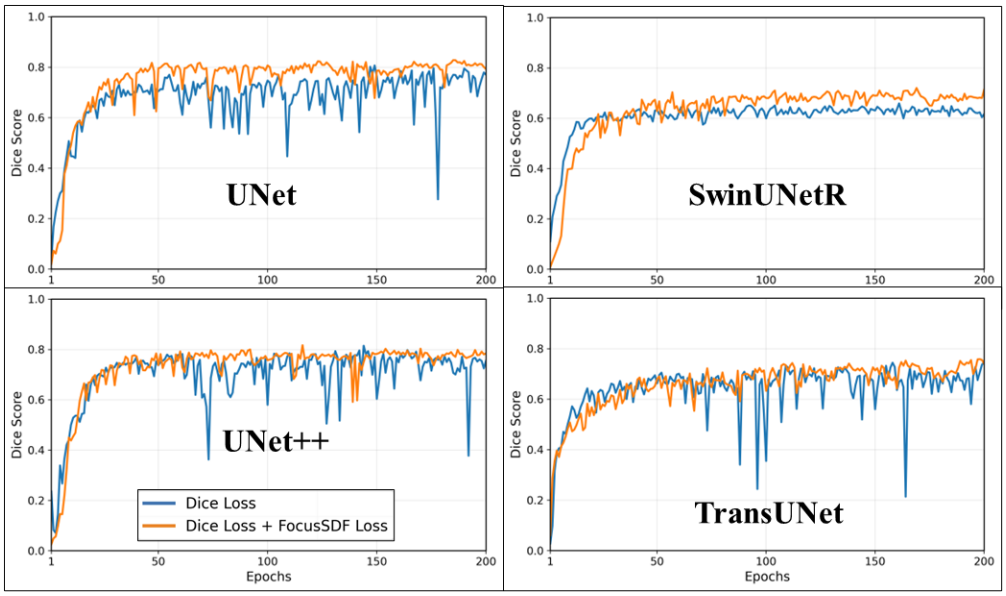

Figure 5.4: Validation performance of UNet, UNet++, SwinUNetR, and TransUNet on the cerebral aneurysm segmentation dataset across training epochs, comparing FocusSDF against competing loss functions. FocusSDF consistently yields smoother validation curves and reduced performance variance, indicating more stable optimization and superior convergence behavior particularly in the boundary-sensitive aneurysm segmentation task.

liver, breast tumor ultrasound, and ischemic stroke MRI datasets yielded Dice scores of 0.8941, 0.7538, and 0.4322, respectively, reflecting better transferability from modalities that were well-represented during pretraining.

Table 5.3: Comparison of FocusSDF against competing loss functions across four datasets and four backbone architectures. UNet, UNet++, SwinUNetR, and TransUNet are evaluated using Dice ↑, IoU ↑, and HD95 ↓ (in pixels). Bold values indicate the best result per column per dataset.

Dataset	**Loss**	**UNet**			**UNet++**			**SwinUNetR**			**TransUNet**		
		Dice↑	**IoU↑**	**HD95↓**	**Dice↑**	**IoU↑**	**HD95↓**	**Dice↑**	**IoU↑**	**HD95↓**	**Dice↑**	**IoU↑**	**HD95↓**
Aneurysm	$\mathcal{L}_{Dice}$	0.535	0.501	37.0	0.433	0.343	39.0	0.554	0.468	32.2	0.545	0.467	30.9
	$\mathcal{L}_1$	0.535	0.501	35.0	0.480	0.447	37.0	0.554	0.468	34.8	0.438	0.407	36.5
	$\mathcal{L}_2$	0.553	0.505	35.5	0.485	0.455	34.0	0.535	0.453	35.2	0.588	0.477	34.9
	$\mathcal{L}_3$	0.496	0.462	36.0	0.548	0.487	34.0	0.422	0.385	35.2	0.299	0.297	38.7
	$\mathcal{L}_4$	0.570	0.528	32.0	0.488	0.418	31.6	0.442	0.328	33.2	0.400	0.383	36.8
	Ours	**0.770**	**0.704**	**25.0**	**0.757**	**0.691**	**23.8**	**0.787**	**0.695**	**22.2**	**0.794**	**0.702**	**21.5**
Liver	$\mathcal{L}_{Dice}$	0.914	0.872	21.2	0.925	0.880	19.0	0.905	0.874	20.4	0.922	0.885	19.1
	$\mathcal{L}_1$	0.920	0.876	21.1	0.914	0.872	19.9	0.872	0.818	22.6	0.907	0.865	17.2
	$\mathcal{L}_2$	**0.951**	**0.925**	**13.8**	0.891	0.833	27.2	0.862	0.792	24.1	0.922	0.885	19.1
	$\mathcal{L}_3$	0.317	0.200	45.4	0.261	0.155	42.5	0.307	0.191	40.3	0.291	0.181	38.0
	$\mathcal{L}_4$	0.898	0.844	22.8	0.903	0.857	20.3	0.902	0.847	17.6	0.927	0.892	14.5
	Ours	0.950	0.915	15.0	**0.950**	**0.916**	**14.7**	**0.934**	**0.875**	**14.8**	**0.969**	**0.936**	**13.8**
Stroke	$\mathcal{L}_{Dice}$	0.705	0.634	26.0	0.712	0.655	25.1	**0.786**	**0.700**	**14.2**	0.755	0.684	24.3
	$\mathcal{L}_1$	0.698	0.611	26.5	0.698	0.567	25.9	0.702	0.597	17.7	0.708	0.599	15.4
	$\mathcal{L}_2$	0.701	0.620	26.2	0.702	0.617	25.7	0.655	0.546	20.3	0.689	0.576	18.4
	$\mathcal{L}_3$	0.346	0.246	48.2	0.365	0.272	46.4	0.325	0.258	41.3	0.206	0.118	38.3
	$\mathcal{L}_4$	0.683	0.602	25.4	0.695	0.606	24.8	0.684	0.575	20.2	0.706	0.597	17.7
	Ours	**0.758**	**0.675**	**20.8**	**0.756**	**0.667**	**18.9**	0.775	0.698	15.5	**0.754**	**0.666**	**15.3**
Breast Tumor	$\mathcal{L}_{Dice}$	0.684	0.635	27.0	0.715	0.642	25.8	0.651	0.584	24.9	0.696	0.629	22.9
	$\mathcal{L}_1$	0.756	0.688	24.6	0.747	0.647	25.1	0.612	0.533	29.6	0.626	0.545	29.7
	$\mathcal{L}_2$	0.747	0.664	22.4	0.758	0.665	24.1	0.606	0.571	28.1	0.499	0.399	30.3
	$\mathcal{L}_3$	0.384	0.310	52.2	0.299	0.226	50.5	0.611	0.181	45.0	0.388	0.329	42.3
	$\mathcal{L}_4$	0.656	0.584	29.7	**0.817**	**0.721**	**18.7**	0.659	0.581	25.4	0.679	0.599	20.6
	Ours	**0.791**	**0.704**	**20.2**	0.802	0.732	19.9	**0.725**	**0.647**	**22.3**	**0.765**	**0.682**	**21.8**

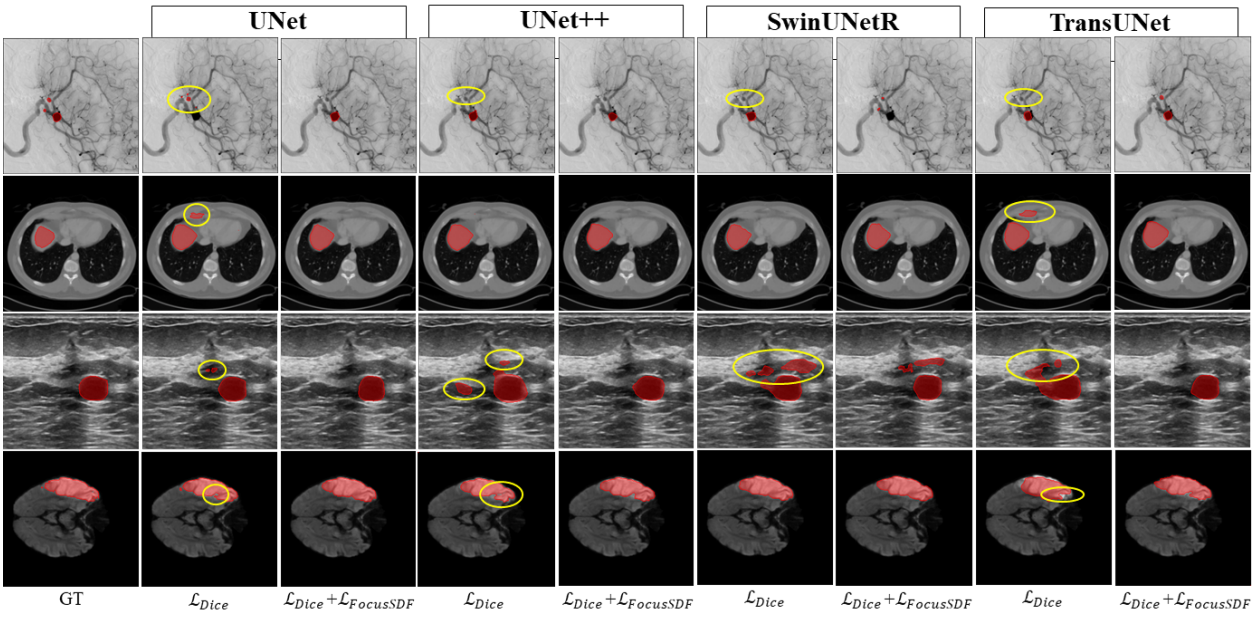

Figure 5.5: Qualitative comparison of segmentation outputs across all four datasets (DSA aneurysm, CTA liver, MRI stroke, and ultrasound breast tumor). Incorporating the proposed FocusSDF loss alongside the Dice loss yields superior boundary precision and structural continuity preservation compared to using the Dice loss alone. Boundary mispredictions produced by Dice loss are highlighted with circles for visual reference.

CHAPTER 6

Conclusion & Future Work

6.1 Conclusion

In this thesis, we presented a unified framework for intracranial aneurysm analysis that integrates structure-aware detection and geometry-aware segmentation within a foundation model-driven paradigm, following the natural clinical imaging pathway from non-invasive screening to high-resolution morphological assessment. The proposed approach addresses two fundamental challenges in neurovascular imaging: anatomically meaningful localization of aneurysms within the cerebrovascular network, and precise boundary delineation of the detected aneurysm sac.

First, we proposed **ARAN**, the first framework for artery-aware intracranial aneurysm localization from CTA and ToF-MRA scans. By integrating a fine-tuned VISTA3D foundation model with vascular geometry descriptors derived from centerline analysis - including vessel radius profiles, cross-sectional eccentricity, and Frenet-Serret curvature and torsion - and fusing them through the proposed GGCA mechanism, ARAN enables structure-aware reasoning for aneurysm detection and parent artery identification. The proposed method consistently outperforms vision-only and geometry-only baselines on both CTA and ToF-MRA datasets, achieving accuracies of 78.2% and 85.8% respectively, demonstrating the effectiveness of incorporating vascular topology and artery-level context for improved detection and clinical interpretability.

Second, we introduced **FocusSDF**, a boundary-aware loss function based on signed distance function supervision, designed to emphasize regions near object boundaries while enforcing gradient consistency in spatial transitions. Once an aneurysm is detected and the clinician orders a DSA scan for detailed morphological evaluation, FocusSDF enables precise delineation of the aneurysm sac with high boundary accuracy. Experimental results across diverse medical imaging datasets - including DSA aneurysm, CTA liver, MRI stroke, and ultrasound breast tumor - demonstrate that FocusSDF consistently improves boundary-sensitive performance, with particularly strong gains in anatomically complex scenarios such as intracranial aneurysm segmentation in DSA. These results highlight the importance of explicit geometric modeling for accurate lesion delineation, especially

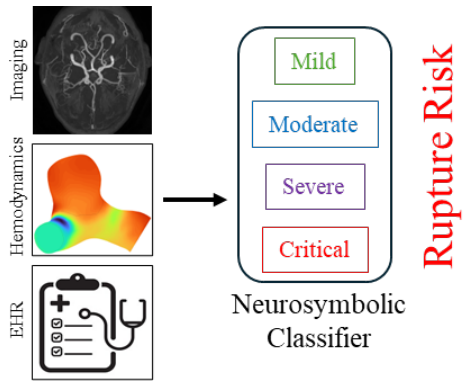

Figure 6.1: Vision for artery-aware rupture risk prediction as the next step beyond this thesis. Given the aneurysm segmentation mask from FocusSDF applied to DSA images, alongside the parent artery context identified by ARAN and complementary clinical metadata from electronic health records (EHR), the envisioned neurosymbolic reasoning system predicts a structured rupture risk level - ranging from Mild to Critical - to directly support neurosurgical decision-making and treatment planning.

when fine structural details are critical for clinical assessment.

Collectively, these contributions establish that **artery-aware detection in CTA and ToF-MRA followed by boundary-aware segmentation in DSA form a clinically coherent and technically robust foundation for advancing toward automated aneurysm rupture risk prediction**. Accurate localization - both in terms of correct identification of the parent artery and precise boundary delineation of the aneurysm sac - is essential for estimating the morphological and hemodynamic factors most strongly associated with rupture risk, and this thesis demonstrates that both can be reliably achieved through complementary foundation model-driven representations.

6.2 Future Work

Building upon the clinically grounded pipeline established in this thesis, future research will focus on developing a **neurosymbolic reasoning-based classifier for artery-aware aneurysm rupture risk prediction**, as illustrated in Fig. 6.1. The envisioned system integrates three complementary streams of evidence - visual morphology, hemodynamic simulation, and electronic health records (EHR) - within a neurosymbolic framework that combines the representational power of deep learning with the interpretability and clinical reasoning capacity of symbolic AI. This moves beyond isolated detection and segmentation toward a holistic, clinically actionable risk stratification system capable of producing structured rupture risk predictions ranging from Mild to Critical.

BIBLIOGRAPHY

[1] Mayo Clinic Staff, "Brain aneurysm: Symptoms and causes," 2025. Mayo Clinic. Accessed: 2026-03-14.

[2] Brain Aneurysm Foundation, "Brain aneurysm statistics & facts," 2023. Accessed: 2026-03-14.

[3] American Heart Association News, "Being aware of an unruptured brain aneurysm may heighten mental health risks," Aug. 2024. Accessed: 2026-03-14.

[4] L. Walkoff, W. Brinjikji, A. Rouchaud, J. Caroff, and D. F. Kallmes, "Comparing magnetic resonance angiography (mra) and computed tomography angiography (cta) with conventional angiography in the detection of distal territory cerebral mycotic and oncotic aneurysms," Interventional Neuroradiology, vol. 22, pp. 524–528, Oct. 2016.

[5] A. Raghuram, R. Patel, A. Varon, R. Sabotin, S. Sanchez, C. P. Derdeyn, P. Jabbour, D. M. Hasan, and E. A. Samaniego, "Volumetric surveillance of brain aneurysms: pitfalls of mra," Interventional Neuroradiology, vol. 29, no. 5, pp. 532–539, 2023.

[6] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al., "Segment anything," in Proceedings of the IEEE/CVF international conference on computer vision, pp. 4015–4026, 2023.

[7] J. Ma, Y. He, F. Li, L. Han, C. You, and B. Wang, "Segment anything in medical images," Nature Communications, vol. 15, p. 654, 2024.

[8] Y. He, P. Guo, Y. Tang, A. Myronenko, V. Nath, Z. Xu, D. Yang, C. Zhao, B. Simon, M. Belue, et al., "Vista3d: A unified segmentation foundation model for 3d medical imaging," in Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 20863–20873, 2025.

[9] 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, no. 1, pp. 25–32, 2019.

[10] A. Park, C. Chute, P. Rajpurkar, J. Lou, R. L. Ball, K. Shpanskaya, R. Jabarkheel, L. H. Kim, E. McKenna, J. Tseng, et al., "Deep learning–assisted diagnosis of cerebral aneurysms using the headxnet model," JAMA network open, vol. 2, no. 6, p. e195600, 2019.

[11] A. Keedy, "An overview of intracranial aneurysms," McGill Journal of Medicine, vol. 9, pp. 141–146, July 2006.

[12] K. M. Malik, S. M. Anjum, H. Soltanian-Zadeh, H. Malik, and G. M. Malik, "A framework for intracranial saccular aneurysm detection and quantification using morphological analysis of cerebral angiograms," IEEE Access, vol. 6, pp. 7970–7986, 2018.

[13] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," arXiv preprint arXiv:1710.10903, 2017.

[14] S. Talukder, "The importance of loss functions in machine learning." https://medium.com, 2025. Accessed: 2026-04-09.

[15] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in International Conference on Medical image computing and computer-assisted intervention, pp. 234–241, Springer, 2015.

[16] Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: A nested u-net architecture for medical image segmentation," in Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support, (Cham), pp. 3–11, Springer International Publishing, 2018.

[17] A. Hatamizadeh, V. Nath, Y. Tang, D. Yang, H. R. Roth, and D. Xu, "Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images," in International MICCAI brainlesion workshop, pp. 272–284, Springer, 2021.

[18] J. Chen, J. Mei, X. Li, Y. Lu, Q. Yu, Q. Wei, X. Luo, Y. Xie, E. Adeli, Y. Wang, M. P. Lungren, S. Zhang, L. Xing, L. Lu, A. Yuille, and Y. Zhou, "Transunet: Rethinking the u-net architecture design for medical image segmentation through the lens of transformers," Medical Image Analysis, vol. 97, p. 103280, 2024.

[19] V. I. Butoi, J. J. G. Ortiz, T. Ma, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "Universeg: Universal medical image segmentation," in Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 21438–21451, 2023.

[20] Y. Du, F. Bai, T. Huang, and B. Zhao, "Segvol: Universal and interactive volumetric medical image segmentation," Advances in Neural Information Processing Systems, vol. 37, pp. 110746–110783, 2024.

[21] J. Ma, Y. He, F. Li, L. Han, C. You, and B. Wang, "Segment anything in medical images," Nature communications, vol. 15, no. 1, p. 654, 2024.

[22] S. Pai, I. Hadzic, D. Bontempi, K. Bressem, B. H. Kann, A. Fedorov, R. H. Mak, and H. J. Aerts, "Vision foundation models for computed tomography," arXiv preprint arXiv:2501.09001, 2025.

[23] Y. Wang, N. Huang, T. Li, Y. Yan, and X. Zhang, "Medformer: A multi-granularity patching transformer for medical time-series classification," Advances in Neural Information Processing Systems, vol. 37, pp. 36314–36341, 2024.

[24] Z. Shi, B. Hu, M. Lu, M. Zhang, H. Yang, B. He, J. Ma, C. Hu, L. Lu, S. Li, et al., "Development and validation of a sham-ai model for intracranial aneurysm detection at ct angiography," Radiology: Artificial Intelligence, vol. 7, no. 3, p. e240140, 2025.

[25] E. Rainville, A. Rasoulian, H. Rivaz, and Y. Xiao, "Weakly supervised intracranial aneurysm detection and segmentation in mr angiography via multi-task unet with vesselness prior," in Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 938–947, 2025.

[26] J. J. Heit, J. M. Honce, V. S. Yedavalli, C. E. Baccin, R. T. Tatit, K. Copeland, and V. M. Timpone, "Rapid aneurysm: Artificial intelligence for unruptured cerebral aneurysm detection on ct angiography," Journal of Stroke and Cerebrovascular Diseases, vol. 31, p. 106690, Oct. 2022.

[27] X. Dai, L. Huang, Y. Qian, S. Xia, W. Chong, J. Liu, A. Di Ieva, X. Hou, and C. Ou, "Deep learning for automated cerebral aneurysm detection on computed tomography images," International Journal of Computer Assisted Radiology and Surgery, vol. 15, no. 4, pp. 715–723, 2020.

[28] S. Aoki, Y. Sasaki, T. Machida, T. Ohkubo, and M. Minami, "Cerebral aneurysms: detection and delineation using 3-d-ct angiography.," American journal of neuroradiology, vol. 13, no. 4, pp. 1115–1120, 1992.

[29] J. N. Stember, P. Chang, D. M. Stember, M. Liu, J. Grinband, C. G. Filippi, P. Meyers, and S. Jambawalikar, "Convolutional neural networks for the detection and measurement of cerebral aneurysms on magnetic resonance angiography," Journal of digital imaging, vol. 32, no. 5, pp. 808–815, 2019.

[30] T. Nakao, S. Hanaoka, Y. Nomura, I. Sato, M. Nemoto, S. Miki, E. Maeda, T. Yoshikawa, N. Hayashi, and O. Abe, "Deep neural network-based computer-assisted detection of cerebral aneurysms in mr angiography," Journal of Magnetic Resonance Imaging, vol. 47, no. 4, pp. 948–953, 2018.

[31] A. M. Ceballos-Arroyo, H. T. Nguyen, F. Zhu, S. M. Yadav, J. Kim, L. Qin, G. Young, and H. Jiang, " Vessel-aware aneurysm detection using multi-scale deformable 3D attention ," in proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, vol. LNCS 15005, Springer Nature Switzerland, October 2024.

[32] S. S. M. Salehi, D. Erdogmus, and A. Gholipour, "Tversky loss function for image segmentation using 3d fully convolutional deep networks," in Machine Learning in Medical Imaging, (Cham), pp. 379–387, Springer International Publishing, 2017.

[33] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in Proceedings of the IEEE international conference on computer vision, pp. 2980–2988, 2017.

[34] J. Ma, Z. Wei, Y. Zhang, Y. Wang, R. Lv, C. Zhu, C. Gaoxiang, J. Liu, C. Peng, L. Wang, Y. Wang, and J. Chen, "How distance transform maps boost segmentation cnns: An empirical

study," in Proceedings of the Third Conference on Medical Imaging with Deep Learning, vol. 121 of Proceedings of Machine Learning Research, pp. 479–492, PMLR, 06–08 Jul 2020.

[35] D. Karimi and S. E. Salcudean, "Reducing the hausdorff distance in medical image segmentation with convolutional neural networks," IEEE Transactions on medical imaging, vol. 39, no. 2, pp. 499–513, 2019.

[36] H. Kervadec, J. Bouchtiba, C. Desrosiers, E. Granger, J. Dolz, and I. Ben Ayed, "Boundary loss for highly unbalanced segmentation," Medical Image Analysis, vol. 67, p. 101851, 2021.

[37] Y. Xue, H. Tang, Z. Qiao, G. Gong, Y. Yin, Z. Qian, C. Huang, W. Fan, and X. Huang, "Shape-aware organ segmentation by predicting signed distance maps," in Proceedings of the AAAI conference on artificial intelligence, vol. 34, pp. 12565–12572, 2020.

[38] F. Gül and E. Aptoula, "A distance transform based loss function for the semantic segmentation of very high resolution remote sensing images," in IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium, pp. 9888–9891, 2024.

[39] A. Myronenko, "3d mri brain tumor segmentation using autoencoder regularization," in International MICCAI brainlesion workshop, pp. 311–320, Springer, 2018.

[40] W. Silversmith, J. A. Bae, P. H. Li, and A. M. Wilson, "Kimimaro: Skeletonize densely labeled 3d image segmentations," 2021. Released: 2021-09-29.

[41] K. Yang, F. Musio, Y. Ma, N. Juchler, J. C. Paetzold, R. Al-Maskari, L. Höher, H. B. Li, I. E. Hamamci, A. 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.

[42] T. D. Noto, G. Marie, S. Tourbier, Y. Alemán-Gómez, O. Esteban, G. Saliou, M. B. Cuadra, P. Hagmann, and J. Richiardi, ""lausanne tof-mra aneurysm cohort"," 2022.

[43] C. M. de Nys, E. S. Liang, M. Prior, M. A. Woodruff, J. I. Novak, A. R. Murphy, Z. Li, C. D. Winter, and M. C. Allenby, ""royal brisbane tof-mra intracranial aneurysm$_{d}atabase$",″ 2024.

[44] Bo et al., "Toward human intervention-free clinical diagnosis of intracranial aneurysm via deep neural network," Elsevier, 2021.

[45] A. E. Kavur, N. S. Gezer, M. Barış, S. Aslan, P.-H. Conze, V. Groza, D. D. Pham, S. Chatterjee, P. Ernst, S. Özkan, B. Baydar, D. Lachinov, S. Han, J. Pauli, F. Isensee, M. Perkonigg, R. Sathish, R. Rajan, D. Sheet, G. Dovletov, O. Speck, A. Nürnberger, K. H. Maier-Hein, G. Bozdağı Akar, G. Ünal, O. Dicle, and M. A. Selver, "Chaos challenge - combined (ct-mr) healthy abdominal organ segmentation," Medical Image Analysis, vol. 69, p. 101950, 2021.

[46] M. R. Hernandez Petzsche, E. De La Rosa, U. Hanning, R. Wiest, W. Valenzuela, M. Reyes, M. Meyer, S.-L. Liew, F. Kofler, I. Ezhov, et al., "Isles 2022: A multi-center magnetic resonance imaging stroke lesion segmentation dataset," Scientific data, vol. 9, no. 1, p. 762, 2022.

[47] A. Shah, "Breast ultrasound images dataset." Download from Kaggle, 2017.

[48] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., "Pytorch: An imperative style, high-performance deep learning library," Advances in neural information processing systems, vol. 32, 2019.

[49] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778, 2016.

[50] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4700–4708, 2017.