



NATIONAL TECHNICAL UNIVERSITY OF ATHENS
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING
SCHOOL OF MECHANICAL ENGINEERING
INTERDISCIPLINARY POSTGRADUATE PROGRAMME
TRANSLATIONAL ENGINEERING IN HEALTH AND MEDICINE

Deep multimodal fusion of image and non-image data in identification of high-risk carotid atheromatous plaque

POSTGRADUATE DIPLOMA THESIS

EVANGELOS STAMOS

Supervisor: Konstantina S. Nikita
Professor, NTUA

Athens, June 2025



NATIONAL TECHNICAL UNIVERSITY OF ATHENS
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING
SCHOOL OF MECHANICAL ENGINEERING
INTERDISCIPLINARY POSTGRADUATE PROGRAMME
TRANSLATIONAL ENGINEERING IN HEALTH AND MEDICINE

Deep multimodal fusion of image and non-image data in identification of high-risk carotid atheromatous plaque

POSTGRADUATE DIPLOMA THESIS

EVANGELOS STAMOS

Supervisor: Konstantina S. Nikita
Professor, NTUA

The postgraduate diploma thesis has been approved by the examination committee on
July 4th, 2025.

(Signature)

(Signature)

(Signature)

.....
Konstantina S. Nikita
Professor, NTUA

.....
Giorgos Stamou
Professor, NTUA

.....
Athanasios Voulodimos
Assistant Professor, NTUA

Athens, June 2025

(Signature)

.....

Evangelos Stamos

Graduate of the Interdisciplinary Postgraduate Programme,
Translational Engineering in Health and Medicine,
Master of Science,
School of Electrical and Computer Engineering,
National Technical University of Athens

Copyright © Evangelos Stamos, 2025 - All rights reserved.

It is prohibited to copy, store and distribute this work, in whole or in part, for commercial or for-profit purposes. It is permissible to reprint, store and distribute for non-profit, educational or research purposes, provided that the source is mentioned and the present message is maintained.

The views and conclusions contained in this work express the author and it should not be construed that they represent official positions of the National Technical University of Athens, including its Schools, Departments, and Units, neither the Supervisor, or the committee that approved it.

Abstract

Cardiovascular diseases remain the leading cause of global mortality, accounting for over 19 million deaths annually according to the World Health Organization. Among these, carotid atherosclerosis, a pathological process characterized by plaque accumulation in the carotid arteries, is a major contributor to ischemic stroke, driven by plaque rupture and thromboembolic events. Traditional diagnostic approaches, which rely on single-modality imaging or clinical data, often fail to capture the complex interplay of morphological, biological, and hemodynamic factors that determine plaque vulnerability. This dissertation addresses this critical gap by developing advanced deep multimodal fusion frameworks to integrate heterogeneous data sources, enabling precise risk stratification of carotid atheromatous plaques.

The primary objective of this work is to create an end-to-end trainable system that synergistically combines B-mode carotid ultrasound imaging with non-image clinical data, including biochemical markers, protein biomarkers, and patient demographics. Three fusion strategies were rigorously investigated: (1) Joint Attention-Based Fusion, which dynamically weights imaging and tabular data contributions through learned attention mechanisms; (2) Early Fusion, merging raw inputs at the feature level; and (3) Attention-Gated Video Hybrid Fusion, a novel architecture designed to process spatiotemporal ultrasound frame sequences alongside clinical data. These models were trained and validated on a multimodal dataset comprising 96 DICOM ultrasound recordings and 73 curated clinical profiles from a cohort of 73 patients, stratified into high-risk (symptomatic with $\geq 50\%$ stenosis or asymptomatic with $\geq 70\%$ stenosis) and low-risk groups.

The Joint Attention-Based Fusion model with an EfficientNet-B0 backbone achieved superior performance, yielding AUC: 86.07%, Balanced Accuracy: 73.28%, F1 Score: 78.42 %, Sensitivity: 81.67%, outperforming unimodal approaches (imaging-only AUC: 84.55%; tabular-only AUC: 64.61%). Attention weights highlighted the dominance of imaging data (69.4% contribution), while clinical biomarkers provided complementary risk context. The proposed Attention-Gated Video Hybrid Fusion framework demonstrated feasibility for dynamic plaque analysis but faced computational constraints in scaling 3D spatiotemporal convolutions.

In conclusion, this thesis advances the field of precision vascular medicine by demonstrating that multimodal deep learning can substantially improve the accuracy and reliability of carotid atheromatous plaque risk assessment compared to conventional single-modality approaches. The proposed architectures, particularly the attention-based fusion paradigm, provide a scalable framework for integrating heterogeneous data sources. Additionally, a prototype Clinical Decision Support System (CDSS) was developed as a web-based interface, illustrating the translational potential of these models for future integration into clinical workflows. Together, these contributions highlight the promise of AI-driven methodologies in enhancing stroke prevention strategies and supporting individualized therapeutic decision-making.

Keywords: Carotid Atherosclerosis, Plaque Vulnerability, Deep Learning, Multimodal Fusion, Stroke Prevention

Contents

Abstract	7
1 Introduction	15
1.1 Background and Context	15
1.2 Problem Statement	16
1.3 Research Objectives	18
1.4 Thesis Structure	18
2 Carotid Atherosclerosis	19
2.1 Carotid Arteries	19
2.2 Atherogenesis	22
2.3 Clinical Implications	24
2.4 Risk Factors	25
2.5 Pathobiological Significance of Circulating Biomarkers	27
2.6 Imaging Modalities for Carotid Plaque Detection	29
2.7 Therapeutic management	31
3 Deep Learning	39
3.1 Deep Learning in Medical Imaging	40
3.2 Data Constraints in Deep Learning for Healthcare	41
3.3 Handling imbalanced dataset strategies	42
3.3.1 Data-Level Approaches	42
3.3.2 Algorithm-Level Approaches	43
3.4 Evaluation Considerations	44
3.5 Applications on Carotid Atherosclerosis	46
4 Multimodal Learning	51
4.1 Data Fusion Strategies	51
4.2 Applications on Biomedical Data	55
5 Dataset	59
5.1 Imaging Data	60
5.2 Tabular Data	64

6 Implementation	73
6.1 Tabular data preprocessing	74
6.2 Imaging data preprocessing	75
6.3 Training and model evaluation	76
6.4 Models Architecture	78
6.4.1 Joint Attention Based Fusion	78
6.4.2 Early fusion	83
6.4.3 Attention-Gated Video Hybrid Fusion	84
6.5 A prototype CDSS for plaque vulnerability assessment	85
7 Results	89
7.1 Joint Attention Based Fusion Models	90
7.2 Ablation study	90
7.3 Early Fusion Model	92
7.4 Summary of Results	93
8 Conclusions and Future Work	97
8.1 Conclusions	97
8.2 Future work	99
Appendices	101
A Supplementary Tables	103
A.1 ESVS Clinical Practice Guidelines on the Management of Atherosclerotic Carotid and Vertebral Artery Disease	103
A.2 AHA atherosclerotic plaque classification	104
Bibliography	113
List of Abbreviations	115

List of Tables

2.1	Carotid Artery Stenosis calculation algorithms	38
2.2	Summary of Plaque-RADS categorization	38
5.1	Stratification scheme followed on patients cohort	59
5.2	Ultrasound Device Settings for Data Collection	61
5.3	Comparison of Biomarker Levels in High-Risk vs. Low-Risk Carotid Plaques	64
5.4	Description of Clinical Parameters Included in the Dataset	69
5.5	Description of Protein Markers Included in the Dataset	70
5.6	Description of Biochemical Parameters Included in the Dataset	71
6.1	One-Hot Encoding of the 'Smoker' Feature	74
6.2	Summary of Tabular Data Features Used as Input in the Model	74
6.3	Detailed Description of Clinical, Biochemical, and Protein Biomarkers Used as Input in the Implemented Models	87
6.4	Image Augmentation Parameters	88
7.1	Summary of set training parameters	89
7.2	Performance Metrics of Joint Fusion Attention-Based Models	90
7.3	Performance Metrics of the Joint Fusion Attention Model with ResNet-18 Backbone at Varying Image Attention Weight Thresholds	91
7.4	Performance Metrics of the Joint Fusion Attention Model with VGG-16 Back- bone at Varying Image Attention Weight Thresholds	91
7.5	Performance Metrics of the Joint Fusion Attention Model with EfficientNet Backbone at Varying Image Attention Weight Thresholds	92
7.6	Ablation Study Findings	92
7.7	Performance Metrics of the Early Fusion Model with Two EfficientNets-B0 .	93
A.1	ESC Classification of Recommendations	103
A.2	ESC Levels of Evidence	103
A.3	AHA Classification of Atherosclerotic Plaque and Its MRI-Modified Counter- part	104

List of Figures

2.1	LCCA stenosis due to atherosclerotic plaque formation	20
2.2	CCA origin	21
2.3	RCCA anatomy	21
2.4	Layers of artery wall - view of transverse section	22
2.5	Schematic representation of atheroma plaque formation	23
2.6	Significant factors associated with carotid atherosclerosis	27
2.7	Carotid DUS - information provided by the Color Doppler study	30
2.8	CTA showing right severe carotid bulb stenosis with deep plaque ulceration	30
2.9	MRA standardized protocol in patients with potential carotid stenting indi- cations	31
2.10	Increased uptake of 18F-FDG (arrows) in LICA plaque on PET/CT scan . .	32
2.11	Illustration of LCEA procedure	33
2.12	CEA - Plaque and shunt	33
2.13	Prior and Post CAS procedure illustration	34
2.14	Management of "average risk" patients with asymptomatic and symptomatic carotid stenoses with BMT, CEA, and/or CAS	35
2.15	Carotid stenosis measurement loci	36
2.16	Plaque-RADS Flowchart	37
2.17	Schematic representation of Plaque-RADS categories	37
3.1	Oversampling and Undersample schematic	42
3.2	Training set division into balanced subsets	43
3.3	Architecture of CFPNet-M	47
3.4	Architecture of (a) Original CFP module (b) CFP module	48
3.5	A highly discordant carotid plaque in the AtheroRisk Software	48
3.6	ROI (far wall of the artery) detection in CCA US	49
3.7	Deep-architecture designed for the LII and MAI segmentation.	49
4.1	MML fusion strategies	52
4.2	Model architecture for different fusion strategies	54
4.3	Overview of HEALNet architecture	56
4.4	Fusion Architectures for breast imaging and tabular data	56
4.5	HyperFusion Framework	57
5.1	Carotid B-Mode US Sample of Dataset	60
5.2	Annotations by radiologists contained in .mat files	62

5.3	Segmented images of atherosclerotic plaque obtained from the same DICOM US file	62
5.4	Patient Cohort Flowchart	63
5.5	Protein markers violin plots I	66
5.6	Protein markers violin plots II	67
5.7	Protein markers normalized radar plot	68
6.1	Data preprocessing workflow	73
6.2	Original and processed image data samples	76
6.3	Patients records distribution in Low and High Risk classes	77
6.4	ResNet18 Architecture	79
6.5	VGG16 Architecture	79
6.6	EfficientNet-B0 Architecture	80
6.7	Joint Attention Based Fusion model architecture	82
6.8	Early Fusion model architecture	83
6.9	Architecture of Attention-Gated Video Hybrid Fusion	84
6.10	CDSS prototype for the assessment of carotid plaque vulnerability	85
6.11	CDSS architecture	86
7.1	Confusion matrices of Joint Fusion Attention model with EfficientNet backbone with no threshold applied	95
7.2	Metrics graph of Joint Fusion Attention model with EfficientNet backbone with no threshold applied across folds	96
7.3	Schematic of Ablation Study	96

Chapter **1**

Introduction

1.1 Background and Context

Cardiovascular diseases (CVDs) are the leading cause of mortality worldwide. In 2021, CVDs accounted for an estimated 19.1 million deaths, as reported by the World Health Organization (WHO). These conditions encompass a broad spectrum of disorders, including coronary artery disease, cerebrovascular disease, peripheral arterial disease, and rheumatic heart disease. Collectively, they contribute to over 85% of deaths associated with myocardial infarction and stroke. Notably, approximately 32% of all global deaths are attributable to cardiovascular disease, a figure that rises to 38% in individuals under the age of 70. The age-adjusted mortality rate for CVD per 100,000 population stands at 239.8, while the age-adjusted prevalence rate is 7,354.1 per 100,000 individuals [1].

The European region presents particularly concerning figures, facing one of the highest levels of CVD mortality globally, with 350–400 strokes occurring per 100,000 population annually, while Greece faces an even more challenging situation with approximately 500 cases per 100,000 individuals [2].

Acute cardiovascular events, such as myocardial infarctions and cerebrovascular accidents (commonly referred to as heart attacks and strokes), predominantly arise from obstructions that disrupt the perfusion of blood to critical organs, specifically the heart and brain. The principal pathological mechanism underlying these obstructions is atherosclerosis, a chronic condition characterized by the progressive accumulation of lipid deposits within the intimal layer of arterial walls supplying these regions. In the context of cerebrovascular accidents, additional etiologies include intracerebral hemorrhage resulting from vascular rupture or ischemic events precipitated by thromboembolic occlusions within cerebral vasculature [3, 4].

The increasing prevalence of CVDs is driven by both demographic transitions and modifiable risk factors. The aging global population and rising life expectancy have expanded the cohort at risk of developing cardiovascular conditions. However, behavioral and environmental determinants, including poor dietary habits, sedentary lifestyles, tobacco consumption, excessive alcohol intake, and air pollution, play a critical role in exacerbating the burden of CVDs. These behavioral risks often culminate in intermediate physiological manifestations such as hypertension, hyperglycemia, dyslipidemia, and

obesity, which significantly contribute to adverse cardiovascular outcomes.

The socioeconomic determinants of cardiovascular health are equally pertinent. Lower socioeconomic status has been strongly correlated with increased incidence and worse outcomes of cardiovascular disease. Recent analyses highlight the disproportionate burden faced by socioeconomically disadvantaged populations, who frequently encounter barriers to accessing preventive healthcare services, advanced diagnostic modalities, and optimal therapeutic interventions [5]. There exists a correlation between the percentage of current health expenditure (CHE) relative to GDP and CVD death rates. Countries with lower CHE as a percentage of GDP tend to experience higher age-standardised CVD mortality rates [3].

Early detection of cardiovascular disease remains paramount to effective management, emphasizing the need for timely intervention through counseling and pharmacological treatment. Behavioral modifications have been demonstrated to mitigate cardiovascular risk significantly. Smoking cessation, salt reduction, increased fruit and vegetable consumption, regular physical activity, and reduced alcohol intake collectively contribute to lower CVD incidence.

The diagnosis of cardiovascular disease follows a structured, evidence-based approach, beginning with a thorough clinical evaluation that includes an assessment of patient symptoms and risk factors. This initial evaluation is then complemented by laboratory testing and imaging modalities, which together provide detailed insights into the nature and severity of vascular pathology. Advanced imaging techniques play a pivotal role in both the diagnosis and prognosis of vascular diseases. Non-invasive modalities such as ultrasonography (US), computed tomography angiography (CTA), and magnetic resonance angiography (MRA) enable comprehensive evaluations of vascular structure and function. Furthermore, the behavior and composition of atherosclerotic plaques in the carotid arteries have emerged as critical prognostic indicators of systemic atherosclerosis, particularly in relation to coronary circulation [6].

The persistent and evolving challenges posed by cardiovascular diseases necessitate a multifaceted approach that includes addressing modifiable risk factors, socioeconomic disparities, and genetic predispositions. Continued research and public health initiatives are essential to mitigate the global impact of CVDs.

1.2 Problem Statement

Carotid atherosclerosis is characterized by the progressive accumulation of atherosclerotic plaques within the carotid arteries, the principal vessels supplying oxygenated blood to the brain. The underlying pathophysiological mechanisms include endothelial dysfunction, lipid accumulation, infiltration of inflammatory cells, and fibrous tissue proliferation, leading to plaque formation that may result in arterial stenosis or occlusion. Importantly, the risk of cerebrovascular events is not solely dictated by the degree of luminal narrowing but is also influenced by plaque composition and stability. Vulnerable plaques, characterized by a lipid-rich necrotic core, thin fibrous cap, and intraplaque hemorrhage, exhibit an increased propensity for rupture, leading to thromboembolic complications [7].

Epidemiological studies indicate that the prevalence of carotid atherosclerosis escalates with advancing age and is compounded by traditional cardiovascular risk factors. The asymptomatic nature of early-stage carotid artery disease often results in delayed diagnosis, frequently culminating in transient ischemic attacks (TIAs) or ischemic strokes. Consequently, early identification through targeted screening in high-risk populations is imperative for mitigating adverse clinical outcomes.

Conventional diagnostic approaches for carotid atherosclerosis primarily rely on imaging modalities such as carotid US, CTA, and MRA. Carotid US is widely utilized as an initial screening tool due to its non-invasive nature, cost-effectiveness, and accessibility, making it a practical choice for large-scale risk assessment. Additionally, US enables the evaluation of plaque morphology and hemodynamic parameters with real-time imaging capabilities. CTA provides high-resolution three-dimensional reconstructions and facilitates detailed compositional analysis of plaques, while MRA offers superior soft tissue contrast, allowing the identification of high-risk features such as intraplaque hemorrhage, lipid-rich necrotic cores, and active inflammation, hallmarks strongly associated with increased risk of plaque rupture and subsequent ischemic events [8]. However, despite these advancements, existing imaging techniques predominantly focus on morphological attributes, often neglecting the dynamic biological processes that contribute to plaque vulnerability, including inflammation and neovascularization.

Management strategies for carotid atherosclerosis encompass both pharmacological and interventional approaches. Medical management focuses on intensive risk factor modification, incorporating antihypertensive therapy, lipid-lowering agents, antiplatelet medications, smoking cessation, and lifestyle modifications. In cases of significant stenosis or symptomatic disease, revascularization procedures such as carotid endarterectomy (CEA) or carotid artery stenting (CAS) are performed to restore cerebral perfusion and mitigate stroke risk. The choice between CEA and CAS is determined by individualized patient factors, including anatomical considerations, comorbidities, and procedural risk profiles.

Current methodologies for assessing atherosclerotic plaque vulnerability predominantly rely on single-modality approaches, which often fail to capture the intricate interplay of morphological, biological, and clinical factors that collectively determine plaque stability. This reductionist approach limits the accuracy and reliability of risk stratification, thereby constraining the ability to predict cerebrovascular events effectively and personalize therapeutic interventions.

Despite significant advances in diagnostic and therapeutic modalities that have improved the management of carotid artery disease, there remains a critical need for continued research to develop robust and reliable risk assessment frameworks and precision-targeted interventions for the evaluation of the vulnerability of carotid atheromatous plaque.

1.3 Research Objectives

The primary aim of this research is to develop an advanced deep multimodal fusion framework designed to integrate diverse data modalities, specifically imaging data and non-imaging clinical information within a cohesive, end-to-end trainable architecture. By capitalizing on the complementary strengths of these heterogeneous data sources, the proposed framework seeks to overcome the limitations inherent in traditional single-modality assessments.

This research focuses on the stratification of patients with carotid artery disease into two distinct risk categories: (1) high-risk patients, for whom surgical intervention is indicated, and (2) low-risk patients, who are candidates for conservative management. The accurate classification of patients by the proposed system is anticipated to serve as a valuable diagnostic tool for clinicians, potentially reducing unnecessary surgical interventions while optimizing patient outcomes.

At the core of this system is the application of multimodal machine learning techniques. These methods facilitate the integration of both imaging and tabular data modalities, thereby enhancing the precision and reliability of carotid disease risk stratification. Specifically, the framework will incorporate clinical data obtained from blood tests, encompassing biochemical analyses and protein markers alongside imaging data derived from B-mode carotid ultrasound recordings. By synthesizing these disparate datasets, the proposed system aims to leverage their complementary characteristics to improve the overall diagnostic process and clinical decision-making.

1.4 Thesis Structure

This thesis is systematically structured to provide a comprehensive exploration of multimodal deep learning for carotid plaque risk assessment, progressing from theoretical foundations to clinical applications. Chapter 1 introduces the epidemiological and clinical significance of carotid atherosclerosis, framing the research objectives and challenges in current diagnostic paradigms. Chapter 2 reviews the pathophysiological mechanisms of plaque formation, imaging modalities, and therapeutic strategies, establishing the biomedical context for the study. Chapter 3 critically examines deep learning methodologies, emphasizing solutions for data constraints and class imbalance in medical applications. Chapter 4 synthesizes multimodal fusion strategies, highlighting their potential to integrate heterogeneous data sources for improved risk stratification. Chapter 5 details the dataset composition, preprocessing pipelines, and biomarker analyses, underscoring the interplay between imaging and clinical features. Chapter 6 presents the implementation of multimodal fusion models, including architectural designs and training protocols. Chapter 7 evaluates model performance across metrics, interpreting results in the context of clinical utility. Finally, Chapter 8 discusses implications, limitations, and future directions, advocating for dynamic plaque analysis and advanced spatiotemporal modeling. This structure ensures a cohesive narrative, bridging technical innovation with translational medical research.

Chapter 2

Carotid Atherosclerosis

Carotid artery disease, primarily driven by atherosclerosis, is a leading contributor to ischemic stroke, accounting for 15–20% of cases globally [9]. This condition arises from the progressive accumulation of lipid-rich plaques within the carotid arteries (Figure 2.1), which can induce luminal stenosis or destabilize, leading to plaque rupture and subsequent thromboembolic events. In Europe, with a population of approximately 715 million, an estimated 1.4 million strokes occur annually, making stroke one of the most significant public health concerns [10]. Each year, stroke is responsible for approximately 1.1 million deaths across the continent, ranking as the second leading cause of mortality following coronary artery disease (CAD) [10, 11]. Projections indicate that the number of individuals living with the long-term effects of stroke is expected to rise substantially, from 3.7 million in 2015 to 4.6 million by 2035, primarily due to demographic shifts associated with an aging population [12]. The economic burden of stroke is considerable, with European healthcare systems allocating approximately €45 billion annually to stroke-related care in 2015, including indirect costs [12, 11]. A similar trend is observed in the United States, where the total financial impact of stroke was estimated at \$ 49.5 billion (€ 43.9 billion) in 2015–2016, with projections suggesting an increase to \$ 129 billion (€ 114 billion) by 2035 [13].

In this chapter, a review is synthesized to present the current understanding of atherosclerosis pathophysiology, its multifaceted clinical relevance, and the complex interplay of modifiable and non-modifiable risk factors that contribute to its development and progression.

2.1 Carotid Arteries

The carotid arteries are vital conduits for cerebral perfusion, supplying oxygenated blood to the brain, head, and neck. Each individual possesses two common carotid arteries: the left common carotid artery and the right common carotid artery. Under typical anatomical conditions, the right common carotid artery (CCA) arises from the brachiocephalic trunk on the right side and left CCA directly from the aortic arch on the left (Figure 2.2). The CCA ascends through the neck and terminates at the upper border of the thyroid cartilage, where it bifurcates into its terminal branches, internal carotid artery (ICA) and external carotid artery (ECA). In the cervical region, the common

carotid arteries are positioned bilaterally, anterolateral to the neck. They are anatomically associated with the cervical sympathetic chain and the transverse processes of the cervical vertebrae, while being enveloped by the prevertebral muscles and deep cervical fascia.




Figure 2.1: *LCCA stenosis due to atherosclerotic plaque formation*
[14]

The ICA is primarily responsible for cerebral circulation, while the ECA supplies blood to the face and scalp (Figure 2.3). Anatomical variations and pathological conditions affecting these arteries, such as stenosis or dissection, can have profound implications for neurological function and systemic health. Furthermore, carotid arteries are central to several therapeutic interventions, including CEA and CAS, which aim to mitigate ischemic risks associated with atherosclerotic plaques. Their accessibility and critical role in maintaining cerebral homeostasis underscore their importance in vascular surgery and neurovascular research.

The carotid artery wall is composed of three distinct layers (Figure 2.4), each playing a critical role in vascular function and structural integrity. Together, these layers ensure the resilience and functionality of the carotid artery under varying physiological and pathological conditions.

1. **Tunica intima**, the innermost layer, consists of endothelial cells supported by connective tissue and elastin fibers, forming a smooth surface that minimizes friction and facilitates blood flow. This layer is directly exposed to hemodynamic forces and is integral to vascular homeostasis and the pathogenesis of atherosclerosis. The endothelium forms an interface between circulating blood or lymph in the lumen and the rest of the vessel wall.
2. **Tunica media** is predominantly composed of smooth muscle cells and elastic fibers, providing mechanical strength and enabling dynamic regulation of arterial diameter in response to blood pressure fluctuations. This layer is crucial for maintaining




Figure 2.2: CCA origin
[15]





Figure 2.3: RCCA anatomy
[15]

vascular tone and accommodating pulsatile blood flow during systolic and diastolic phases.

3. **Tunica adventitia**, the outermost layer, consists of dense connective tissue interspersed with collagen and elastin fibers, as well as vasa vasorum that supply nutrients to the artery wall. It anchors the artery to surrounding tissues and houses autonomic nerves that regulate smooth muscle contraction.



© 2010 Encyclopædia Britannica, Inc.

Figure 2.4: *Layers of artery wall - view of transverse section*
[16]

2.2 Atherogenesis

The term atherogenesis was first recorded in 1948 by American pathologists G.L. McMillan and G.L. Duff in their study published in the Archives of Pathology. Atherogenesis, from the Greek ἀθήρα meaning "gruel" or "porridge" and γένεσις meaning "origin or formation," is the insidious process of atherosclerotic plaque formation that begins early in life, often decades before clinical manifestations become apparent. This complex pathophysiological cascade initiates with endothelial dysfunction, triggered by various factors such as disturbed wall shear stress (WSS), oxidative stress, and systemic risk factors, leading to the gradual luminal narrowing of the artery and the activation of inflammatory pathways. Atheromatous plaques tend to develop preferentially in arterial regions where hemodynamic forces create conditions conducive to endothelial dysfunction. They are more likely to form in areas of disturbed flow dynamics, such as bifurcations, curvatures,

and regions of low WSS. The carotid bifurcation, the aortic arch, and the ostia of major branches represent key sites of plaque accumulation due to the presence of complex flow patterns, which foster endothelial activation and local inflammatory responses. These regions experience oscillatory and turbulent flow, reducing the protective effects of laminar shear stress.

The process of atherogenesis initiates as the endothelium becomes increasingly permeable. Low-density lipoprotein (LDL) particles infiltrate the arterial intima, where they undergo oxidative modification. These oxidized LDL (oxLDL) molecules act as damage-associated molecular patterns (DAMPs), further exacerbating endothelial damage and igniting a maladaptive inflammatory response (Figure 2.5). Concurrently, circulating monocytes adhere to the activated endothelium, migrate to the sub-endothelial space, and differentiate into macrophages. These macrophages, upon ingesting oxLDL, transform into lipid-laden foam cells, forming the hallmark fatty streaks of early atherosclerosis. As the process advances, vascular smooth muscle cells (VSMCs) proliferate and migrate from the tunica media to the intima, synthesizing extracellular matrix components that contribute to the fibrous cap formation. This intricate interplay between lipid accumulation, inflammation, and vascular remodeling ultimately culminates in the development of advanced atherosclerotic plaques, characterized by a necrotic core, calcification, and potential instability.




Figure 2.5: *Schematic representation of atheroma plaque formation*
[17]

Atheroma progression exhibits diverse patterns, reflecting the complexity of its underlying mechanisms. One common trajectory is the gradual growth of plaques, characterized by the slow accumulation of lipids, inflammatory cells, and extracellular matrix over many years. This process often leads to progressive narrowing of the arterial lumen, potentially resulting in clinical symptoms such as angina when blood flow becomes significantly obstructed. Elevated levels of LDL cholesterol and systolic blood pressure (SBP) have been

identified as key accelerators of this gradual progression, particularly in younger populations [18]. Conversely, some patients experience a much more rapid progression of atherosclerosis, occurring over a period as short as a few months to 2–3 years, even in the absence of conventional risk factors for accelerated disease [19]. Factors contributing to this rapid progression include coronary vasospasm, human immunodeficiency virus (HIV) infection, elevated inflammatory markers, and specific genetic mutations. Additionally, accelerated atherosclerosis is frequently observed in clinical contexts involving significant endothelial injury, such as heart transplantation, coronary artery bypass grafting (CABG), and percutaneous transluminal coronary angioplasty (PTCA).

Intraplaque hemorrhage (IPH) is a key phenomenon in atheroma progression and is a hallmark of vulnerable plaques. IPH has been strongly associated with an increased risk of stroke and is linked to both greater arterial narrowing (lumen stenosis) and accelerated plaque growth. The underlying mechanism is thought to involve the rupture of fragile neo-vessels within the plaque, driven by localized factors such as hemodynamic stress from blood flow, calcification within the plaque structure, and heightened inflammation. Systemic factors also play a critical role; for instance, pulse pressure has been identified as an independent predictor of IPH, alongside hypertension and male sex. IPH not only promotes immediate plaque progression but also drives long-term changes within the vessel wall. This process is accompanied by elevated levels of inflammatory cytokines within the plaque microenvironment. Several risk factors have been linked to IPH progression, including baseline use of antiplatelet agents, a history of smoking, and larger initial carotid plaque volumes.

2.3 Clinical Implications

The clinical manifestation of carotid disease can be symptomatic and asymptomatic. A patient is classified as symptomatic if they have experienced at least one episode of stroke or TIA within the past six months. In the absence of such events, the disease is considered asymptomatic and is most often detected incidentally during routine medical evaluations or imaging for unrelated conditions. An individual may have carotid artery disease for many years without exhibiting symptoms until the underlying pathophysiological processes result in reduced cerebral blood flow, leading to a stroke or TIA. Due to its asymptomatic nature in the early stages, carotid artery disease is often referred to as a "silent killer", as it can progress unnoticed until a severe event occurs.

A TIA is a temporary episode of stroke-like symptoms resulting from a brief disruption of blood flow to the brain. Although a TIA typically lasts only a few minutes and does not cause permanent damage, it serves as a critical warning sign. The INSIST study from Australia reported a one-year ischemic stroke risk of 3.2% following TIA or minor stroke, though this wasn't specifically broken down by age groups [20]. The risk of stroke after TIA appears to be influenced by age not only directly but also through its association with other risk factors that become more prevalent with advancing age, such as hypertension, diabetes, and atrial fibrillation [21, 22]. Often referred to as a "mini-stroke," a TIA highlights an underlying vascular issue and presents an opportunity for

early intervention to mitigate the risk of a more severe cerebrovascular event. Symptoms of TIA happen suddenly and may include [23]:

- Weakness, numbness or paralysis in the face, arm or leg, typically on one side of the body
- Slurred speech or trouble understanding others
- Blindness in one or both eyes or double vision
- Dizziness or loss of balance or coordination

Carotid artery disease, in its most severe manifestations, can precipitate ischemic stroke through three primary pathophysiological mechanisms [24, 25].

1. **Total occlusion** of the carotid artery due to progressive atherosclerotic plaque accumulation, which may culminate in complete obstruction of cerebral perfusion
2. **Thrombosis**, wherein an atherosclerotic plaque undergoes rupture, compromising the integrity of the endothelial lining of the artery resulting in the formation of an intraluminal thrombus that can further impede arterial flow.
3. **Embolism**, which accounts for the majority of stroke events. In this process, a thrombus or a fragment of disrupted plaque detaches from the carotid artery and travels through the bloodstream to cerebral vasculature, where it occludes a distal artery, leading to cerebral ischemia and infarction.

While symptomatic carotid atherosclerosis often manifests as TIAs or strokes, a substantial proportion of high-risk plaques remain clinically silent, posing significant diagnostic and therapeutic challenges. The pathophysiology of carotid atherosclerosis is multifactorial, involving complex interactions between systemic risk factors, such as hypertension, dyslipidemia, diabetes mellitus, smoking, and local vascular mechanisms, including endothelial dysfunction and extracellular matrix remodeling.

Advances in imaging modalities and molecular diagnostics have enhanced the identification of vulnerable plaques; however, the precise mechanisms underlying plaque instability and rupture remain an area of active investigation. Understanding these processes is essential for developing targeted interventions aimed at mitigating stroke risk in affected populations.

2.4 Risk Factors

The pathogenesis of carotid atherosclerosis is deeply intertwined with established cardiovascular risk factors, many of which demonstrate consistent associations across multiple epidemiological studies. Evidence from comprehensive meta-analyses has identified factors significantly associated with carotid plaque formation, including hyperlipidemia,

hyperhomocysteinemia, hypertension, hyperuricemia, smoking, metabolic syndrome, hypertriglyceridemia, diabetes mellitus, and elevated LDL levels [26]. Dyslipidemia constitutes another critical risk category, with hyperlipidemia increasing atherosclerosis risk by greater than 50%. The specific lipid fractions demonstrate differential impacts: elevated LDL directly contributes to plaque formation, while HDL exerts protective effects through promoting reverse cholesterol transport. Hypertension stands as a predominant risk factor, with recent research confirming its independent association with all ischemic stroke subtypes, regardless of carotid plaque burden [27]. The severity of hypertension appears particularly consequential, as grade 3 hypertension has been identified as an independent factor influencing gray-scale median (GSM) values in asymptomatic patients [28]. Diabetes mellitus represents both an independent risk factor and a potent modifier of atherosclerotic progression, with recent data suggesting that fasting plasma glucose levels significantly correlate with GSM values in symptomatic patients [28]. The metabolic dysregulation associated with diabetes accelerates endothelial dysfunction and vascular inflammation, creating a pro-atherogenic environment. Smoking maintains its status as a substantial contributor to carotid atherosclerosis, with current smoking demonstrating stronger associations with disease progression than former smoking.

Recent advances in genetic epidemiology have illuminated the substantial role of genetic predisposition in carotid atherosclerosis susceptibility. Multiple studies have identified specific single nucleotide polymorphisms (SNPs) associated with increased risk in Chinese populations, including IL1A rs1609682 TT and HABP2 rs7923349 TT genotypes [29, 30]. These genetic variants appear to function synergistically, as demonstrated by generalized multifactor dimensionality reduction (GMDR) analysis revealing significant gene-gene interactions among HABP2 rs7923349, ITGA2 rs1991013, IL1A rs1609682, and NOS2A rs8081248. The transcription factor BACH1 (BTB and CNC homology 1) has been implicated in atherosclerotic pathophysiology, with its deletion shown to decrease turbulent blood flow-induced atherosclerotic lesions and macrophage content in plaques [31]. Coronary artery disease-associated risk variant rs2832227 has been linked to BACH1 gene expression in carotid plaques, suggesting a mechanistic connection between genetic predisposition and vascular inflammation. Most recently, a 2024 study revealed that exosomal miR-155-5p derived from periodontal endothelial cells promotes carotid atherosclerosis by increasing endothelial permeability and angiogenic activity, providing a molecular mechanism linking periodontitis to carotid atherosclerosis [32]. Significantly elevated expression of miR-155-5p was detected in plasma exosomes of patients with both chronic periodontitis and carotid atherosclerosis compared to those with periodontitis alone.

Beyond established risk determinants, several emerging factors warrant clinical attention. Hyperhomocysteinemia has been identified as a potent risk factor capable of elevating atherosclerosis risk by more than 50%. The exact pathophysiological mechanisms remain incompletely understood but likely involve endothelial damage, increased oxidative stress, and impaired vasodilation. Metabolic syndrome, representing a constellation of interrelated metabolic abnormalities, significantly associates with carotid plaque presence. This syndrome's contribution likely exceeds the sum of its individual components,

suggesting synergistic interactions among metabolic perturbations. Several behavioral and environmental factors have gained recognition as significant contributors to carotid atherosclerosis. Negative emotions, socioeconomic strain, alcohol consumption, and air pollution exposure have all demonstrated associations with atherosclerotic burden. Obstructive sleep apnea syndrome (OSAS) has emerged as a risk factor through mechanisms involving intermittent hypoxia, sympathetic activation, and systemic inflammation.




Figure 2.6: *Significant factors associated with carotid atherosclerosis*

MetS: metabolic syndrome; Hcy: homocysteine; OSAS: obstructive sleep apnea syndrome; CPAP: continuous positive airway pressure

Risk factors of meta-analysis are above the dotted line.

Risk factors of systematic review are below the dotted line.

The dots with four different filled ratios below risk factors represent different total sample size ranges.

Different colors represent different quality score ranges.

[26]

2.5 Pathobiological Significance of Circulating Biomarkers

In recent years, increasing attention has been directed toward the investigation of systemic circulating biomarkers obtained through peripheral blood sampling, which may offer insights not only into the presence of atherosclerotic disease but also into the characteristics of vulnerable atherosclerotic plaques and the potential risk of future cardiovascular events in affected individuals. The coordinated interaction of systemic and local inflammatory mediators is fundamental to the modulation of the atherogenic pathophysiological continuum; therefore, the identification of circulating biomarkers is an essential component of risk stratification and the development of novel therapeutics. Throughout the progression of atherosclerotic plaque, certain molecules can diffuse into the bloodstream, serving as circulating biomarkers that offer insight into the presence, characteristics, and potential complications of the plaque. For such biomarkers to be regarded as valid surrogates for clinically meaningful cardiovascular outcomes, they must meet

a series of stringent criteria. As outlined by the American Heart Association (AHA), the comprehensive evaluation of a surrogate biomarker for cardiovascular endpoints involves multiple stages: establishing proof of concept, conducting prospective validation, demonstrating incremental value, confirming clinical utility, assessing effects on clinical outcomes, and evaluating cost-effectiveness [33].

C-reactive protein (**CRP**) is recognized as one of the most prominent biomarkers of inflammation, with both standard and high-sensitivity CRP (**hs-CRP**) assays extensively utilized in clinical settings for the stratification of vascular disease. A growing body of evidence has demonstrated associations between elevated hs-CRP levels and the presence of unstable carotid artery stenosis [34], ICA stenosis [35], and identification of individuals at elevated vascular risk [36]. Notably, increased hs-CRP concentrations have been linked to reduced echogenicity of carotid plaques [37], implying a potential relationship between hs-CRP levels and plaque vulnerability. Moreover, vulnerable atherosclerotic plaques have been shown to exhibit upregulated hs-CRP expression. Importantly, elevated hs-CRP levels are independently correlated with a heightened risk of ischemic stroke [38] and are considered among the key risk factors for acute anterior circulation stroke [39].

Pro-inflammatory cytokines such as Interleukin-6 (**IL-6**) and Tumor Necrosis Factor-alpha (**TNF α**) play central roles in endothelial activation, leukocyte recruitment, and the amplification of intraplaque inflammatory cascades. IL-6, a master regulator of the acute-phase inflammatory response, has been identified as an independent predictor of progressive atherosclerosis, with elevated systemic levels conferring an increased propensity for plaque rupture and subsequent cerebrovascular insult [40, 41]. Furthermore, IL-6 has been implicated in the presence of ICA stenosis. Evidence from both clinical observations [35] and genetic association studies supports a link between elevated IL-6 levels and the development of ICA stenosis [42]. Likewise, TNF α , a pro-inflammatory cytokine active in the early stages of the inflammatory response, has been associated with both the prevalence and severity of carotid artery stenosis, as well as with an elevated risk for the development of carotid plaques [43, 44]. Vulnerable atherosclerotic plaques have been shown to exhibit increased expression of TNF α , with higher circulating levels correlating with larger plaque size [45, 42]. Immunohistochemical analyses have further demonstrated that TNF α , in conjunction with hypoxia and oxidized LDL, significantly enhances the expression of matrix metalloproteinase-7 (**MMP-7**), a molecule closely linked to symptomatic carotid artery disease [46]. Moreover, post-CEA plaque assessments revealed significantly elevated TNF α levels in symptomatic individuals [47].

Furthermore, the matrix metalloproteinase (MMP) system, particularly **MMP-9**, and its endogenous tissue inhibitors **TIMP-1** and **TIMP-2**, play crucial roles in extracellular matrix degradation and the preservation of fibrous cap integrity. Dysregulated MMP-TIMP homeostasis has been shown to facilitate the degradation of structural collagen, thereby predisposing plaques to rupture and subsequent thromboembolic complications [48]. These proteolytic imbalances are increasingly recognized as pivotal determinants of plaque instability and cerebrovascular risk.

Among the most clinically and mechanistically salient biomarkers implicated in carotid atherosclerosis is **Galectin-3**, a β -galactoside-binding lectin that serves as a key regulator

of monocyte-macrophage activation, fibrotic remodeling, and extracellular matrix deposition. Elevated plasma concentrations of Galectin-3 have been robustly correlated with increased carotid intima-media thickness (cIMT), a well-established surrogate marker of atherosclerotic burden, as well as heightened plaque vulnerability [49, 50, 51].

Another biomarker of emerging significance is **Retinol-Binding Protein 4 (RBP4)**, an adipokine with integral functions in insulin resistance and lipid metabolism. Empirical evidence suggests that elevated circulating concentrations of RBP4 are independently associated with increased total plaque area and augmented plaque vulnerability, irrespective of conventional cardiovascular risk factors [52].

Collectively, the elucidation of these biomarkers furnishes critical insights into the molecular determinants underpinning the progression and destabilization of atherosclerosis. Their integration into contemporary clinical paradigms holds immense promise for refining predictive algorithms, guiding risk stratification, and informing precision-targeted therapeutic interventions aimed at mitigating the burden of cerebrovascular disease.

2.6 Imaging Modalities for Carotid Plaque Detection

Accurate detection and characterization of carotid plaques are paramount for assessing cerebrovascular risk and guiding clinical decision-making. The primary imaging modalities employed in clinical practice include carotid duplex ultrasound (DUS), CTA and MRA. Each of these techniques offers distinct advantages and limitations with respect to spatial resolution, tissue characterization, and clinical applicability.

Carotid duplex ultrasound (Figure 2.7) is the most widely utilized modality due to its non-invasive nature, cost-effectiveness, and capability to provide real-time hemodynamic assessments. It integrates B-mode imaging, color Doppler, and spectral Doppler techniques to evaluate luminal narrowing and plaque morphology. However, despite its utility, DUS exhibits inherent limitations, including operator dependency, restricted field of view, and suboptimal sensitivity for detecting intraplaque hemorrhage or lipid-rich necrotic cores [53].

CTA (Figure 2.8) provides high-resolution imaging of vascular structures and is particularly adept at delineating calcified plaques and quantifying luminal stenosis. The technique involves intravenous administration of iodinated contrast agents and utilizes multislice computed tomography to generate three-dimensional reconstructions of the carotid arteries. A key advantage of CTA is its ability to visualize both luminal and extraluminal plaque components, thus facilitating risk stratification. However, exposure to ionizing radiation and potential nephrotoxicity from contrast media remain significant concerns, particularly in patients with renal impairment or iodine allergies [53].

MRA (Figure 2.9) serves as a powerful alternative for evaluating carotid atherosclerosis, offering high soft-tissue contrast without ionizing radiation. Utilizing various pulse sequences such as time-of-flight (TOF) and contrast-enhanced MRA, this modality enables precise characterization of plaque composition, including the identification of lipid cores, fibrous caps, and intraplaque hemorrhage. Moreover, the advent of black-blood MRI techniques has further enhanced the detection of high-risk plaque features, such as




Figure 2.7: *Carotid DUS - information provided by the Color Doppler study*
[54]




Figure 2.8: *CTA showing right severe carotid bulb stenosis with deep plaque ulceration*
A, Axial source image. B, Sagittal reformat. C, 3D rendered image (large arrow, stenotic ICA; dashed arrow, bulb plaque ulcer; small arrow, proximal ECA)
[55]

neovascularization and inflammation. Despite these advantages, MRA is constrained by longer acquisition times, motion artifacts, and contraindications in patients with certain metallic implants or gadolinium-based contrast agent allergies [56, 53].

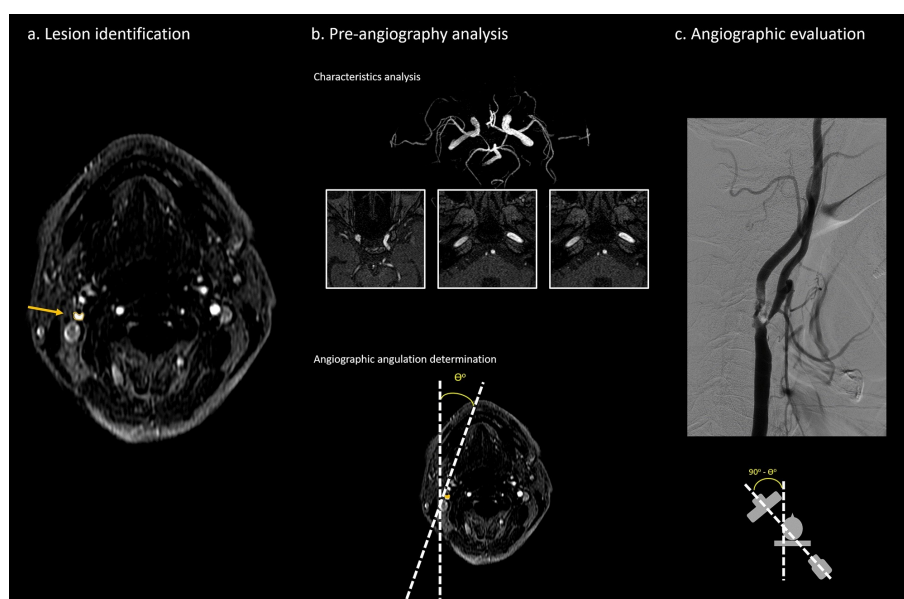


Figure 2.9: *MRA standardized protocol in patients with potential carotid stenting indications*

A: ICA stenosis identified on TOF-MRA; B: The three imaging characteristics, including differences in the opacification grades of bilateral ophthalmic arteries and the signal intensities and diameters of bilateral petrous ICA, were evaluated. An angulation for further DSA was calculated to avoid discrepancy between operators; C: Digital Subtractive Angiography (DSA) was performed based on the previously acquired angulation.

[57]

Emerging imaging techniques such as high-resolution vessel wall imaging, positron emission tomography (PET)-MRI hybrid imaging, PET-CT (Figure 2.10) and advanced ultrasound elastography are under investigation for their potential to improve plaque vulnerability assessment. These modalities aim to enhance risk prediction beyond traditional luminal stenosis quantification by providing insights into plaque biology, inflammation, and mechanical properties [53].

Each imaging technique plays a critical role in modern diagnostic workflows, with the choice of modality contingent on clinical indications, patient-specific factors, and institutional resources. A multimodal imaging approach, integrating complementary strengths of these techniques, may ultimately offer the most comprehensive assessment of carotid atherosclerosis and stroke risk stratification.

2.7 Therapeutic management

The therapeutic management of carotid disease encompasses both conservative and interventional approaches, with the selection of treatment strategies contingent upon the severity of arterial involvement, symptomatology, and individual patient risk factors. Conservative management, which includes pharmacological therapy and lifestyle modifications, serves as the foundation of care, aiming to mitigate disease progression and




Figure 2.10: Increased uptake of ^{18}F -FDG (arrows) in LICA plaque on PET/CT scan
Suggestive of acute inflammation
A: Coronal plane; B: Axial plane.
[58]

reduce the overall cardiovascular risk burden. In cases where conservative measures are insufficient, particularly in patients with significant stenosis or symptomatic presentations, interventional procedures such as CEA or CAS may be warranted. The distinction between symptomatic and asymptomatic patients is of critical importance, as it informs the therapeutic decision-making process, ensuring that interventions are tailored to optimize clinical outcomes while minimizing procedural risks.

1. **Conservative** management involves antihypertensive therapy aimed at achieving optimal blood pressure control, with target values of $\leq 140/90$ mmHg, except for patients with DM or kidney disease, in whom blood pressures should be $\leq 130/80$ mmHg. Concurrently, lipid-lowering treatment, predominantly utilizing statins, is employed to reduce LDL cholesterol levels to ≤ 100 mg/dL. This approach is often complemented by measures to decrease triglyceride concentrations to <150 mg/dL and elevate HDL cholesterol levels to > 40 mg/dL. In addition to pharmacological interventions, lifestyle modifications such as regular physical activity, smoking cessation, weight reduction, and adherence to a balanced and nutritious diet are integral components of comprehensive cardiovascular risk reduction strategies.

2. **Interventional** treatment

(a) **CEA** - Carotid Endarterectomy

A longitudinal incision is made along the anterior aspect of the neck to provide access to the affected carotid artery. Following meticulous dissection and vascular control, the artery is longitudinally arteriotomized, and the atheromatous plaque is carefully excised to restore luminal patency and optimize cerebral perfusion (Figures 2.11, 2.12). Subsequent to plaque removal, the arteriotomy is closed using primary suturing or augmented with a patch graft, which may be composed of autologous venous tissue or synthetic material, to reduce the risk of restenosis and maintain hemodynamic stability.

Left Carotid Endarterectomy




Figure 2.11: *Illustration of LCEA procedure*
[59]




Figure 2.12: *CEA - Plaque and shunt*
The yellow plaque is visible inside the opened CCA and ICA as well as the inserted shunt for brain vascularization.
[60]

(b) **CAS** - Carotid Artery Stenting

A minimally invasive procedure typically performed under local anesthesia with mild sedation. It involves inserting a catheter (Figure 2.13), usually via the femoral artery, to navigate to the carotid artery under fluoroscopic guidance. An embolic protection device is placed to prevent debris from reaching the brain. A self-expanding stent is then deployed across the stenotic lesion to restore blood flow, often followed by balloon angioplasty to optimize stent expansion. The procedure concludes with the removal of the protection device and catheter, with post-procedural monitoring to ensure stability.




Figure 2.13: *Prior and Post CAS procedure illustration*

A: Prior CAS; B: CAS performed, catheter inserted; C: Post CAS
[61]

The discussion on carotid revascularization, particularly the comparison between CEA and CAS, remains a pertinent subject in vascular surgery. CEA has long been established as a reliable intervention for symptomatic patients with significant stenosis, demonstrating substantial benefits in stroke prevention. In contrast, CAS is increasingly considered as an alternative, especially for patients deemed high risk for open surgical procedures. According to the latest European Society for Vascular Surgery (ESVS) guidelines (Figure 2.14), the choice between these techniques should be tailored to individual patient characteristics, including comorbid conditions, anatomical considerations, and overall risk profile. The chromatic differentiation of the rectangular enclosures corresponds to the strata of evidentiary support within the ESVS clinical practice recommendations, as delineated in Figure A.1 of the Supplementary Materials. The algorithmic framework employs data derived from diagnostic imaging modalities, quantifies the severity of luminal stenosis, and incorporates patient age as critical parameters. The AHA atherosclerotic plaque classification (Table A.3) provides an additional layer of diagnostic detail, helping to characterize plaque morphology, which can further inform the choice of revascularization strategy. It is imperative to acknowledge that the stenosis grading within this algorithm adheres to the North American Symptomatic Carotid Endarterectomy Trial (NASCET)

methodology. The principal algorithms employed for the quantification of carotid artery stenosis are tabulated in Table 2.1, while a schematic representation of the measurement loci is provided in Figure 2.15. The core tenets of the ESVS management framework can be synthesized as follows:

- Patients exhibiting significant carotid stenosis, defined as exceeding 60% in asymptomatic individuals and 50% in symptomatic individuals, are candidates for invasive intervention. CEA is generally favored over stenting in these cases.
- Asymptomatic patients of advanced age, specifically those whose age is within five years of their estimated life expectancy, are typically managed conservatively via medical therapy, rather than undergoing interventional procedures.
- Individuals, regardless of symptomatic presentation, who present with complete arterial occlusion are managed pharmacologically, owing to the elevated risk of perioperative complications associated with interventional approaches.

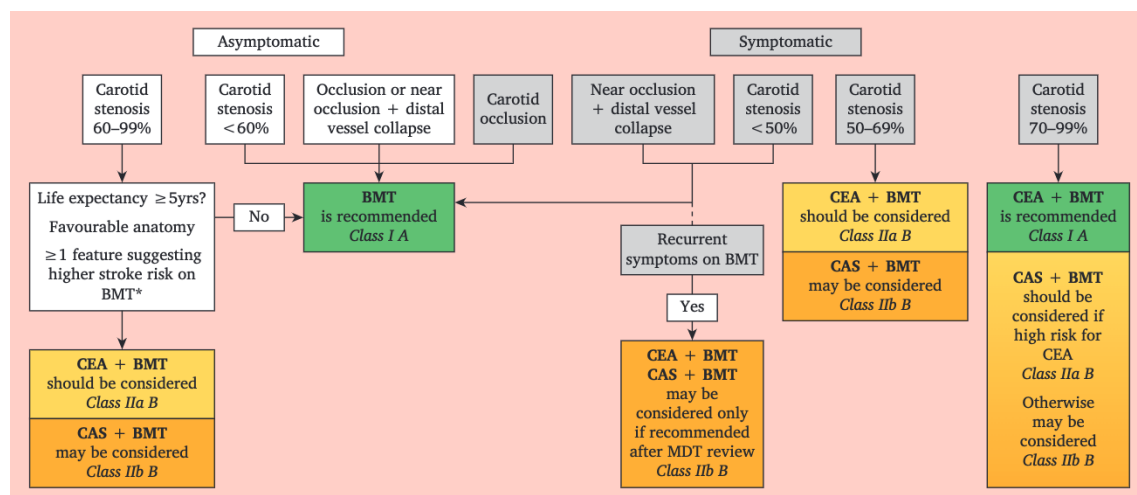


Figure 2.14: Management of "average risk" patients with asymptomatic and symptomatic carotid stenoses with BMT, CEA, and/or CAS

[11]

In 2024, the Carotid Plaque-RADS score was introduced [63], a classification system for carotid plaque and its defining characteristics. Plaque-RADS constitutes a standardized and robust framework for the characterization and reporting of carotid plaque composition and morphology across various imaging modalities, including US, CTA, and MRA. This scoring system holds significant potential for improving the stratification of patients, facilitating the accurate identification of individuals who may benefit from conservative medical management as opposed to those necessitating alternative therapeutic approaches. The adoption of a uniform lexicon and structured reporting format is poised to enhance interdisciplinary communication between radiologists and referring clinicians.

As shown in Figure 2.16, the Plaque-RADS flowchart guides the classification process based on various imaging findings and risk factors, including plaque characteristics such as wall thickness and the presence of a lipid-rich necrotic core (LRNC) or fibrous cap (FC).




Figure 2.15: *Carotid stenosis measurement loci*
[62]

Further, Figure 2.17 provides a schematic representation of the different Plaque-RADS categories, highlighting critical plaque components such as LRNC, maximum wall thickness (MWT), and the integrity of the fibrous cap, all of which influence the risk stratification for cerebrovascular events.

The detailed categorization of plaque types and their corresponding risks is summarized in Table 2.2. As outlined in the table, plaques with a higher Plaque-RADS score, such as scores 4a, 4b and 4c, correlate with increased risk for cerebrovascular events, especially when complicated features such as IPH, ruptured fibrous cap, or intraluminal thrombus are present.

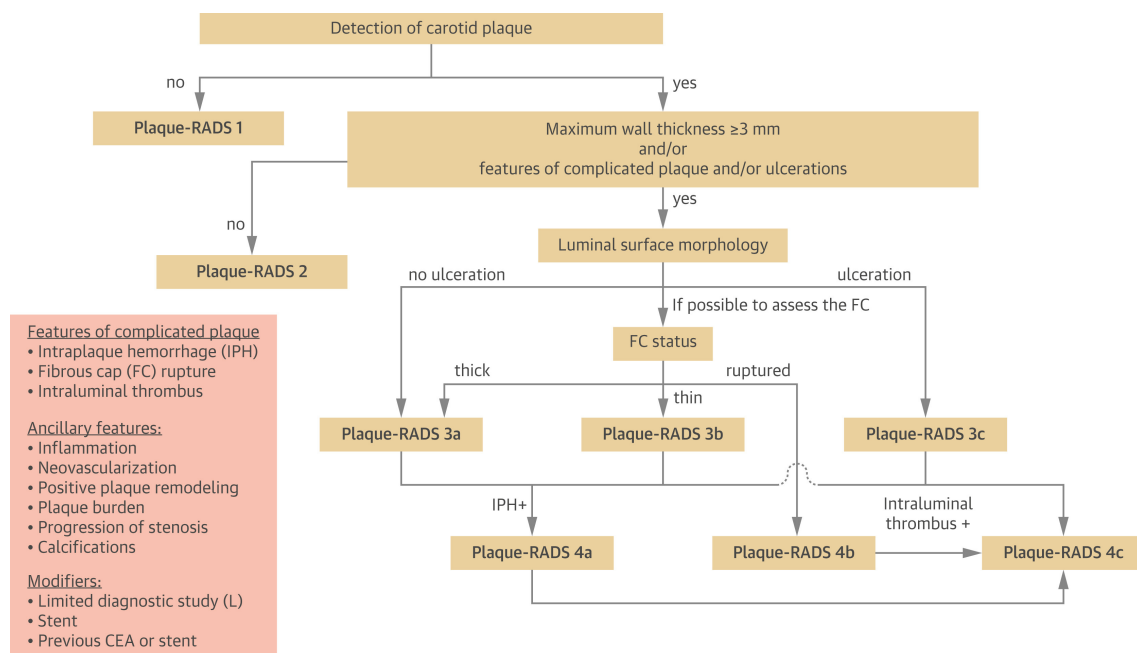


Figure 2.16: *Plaque-RADS Flowchart*
[63]

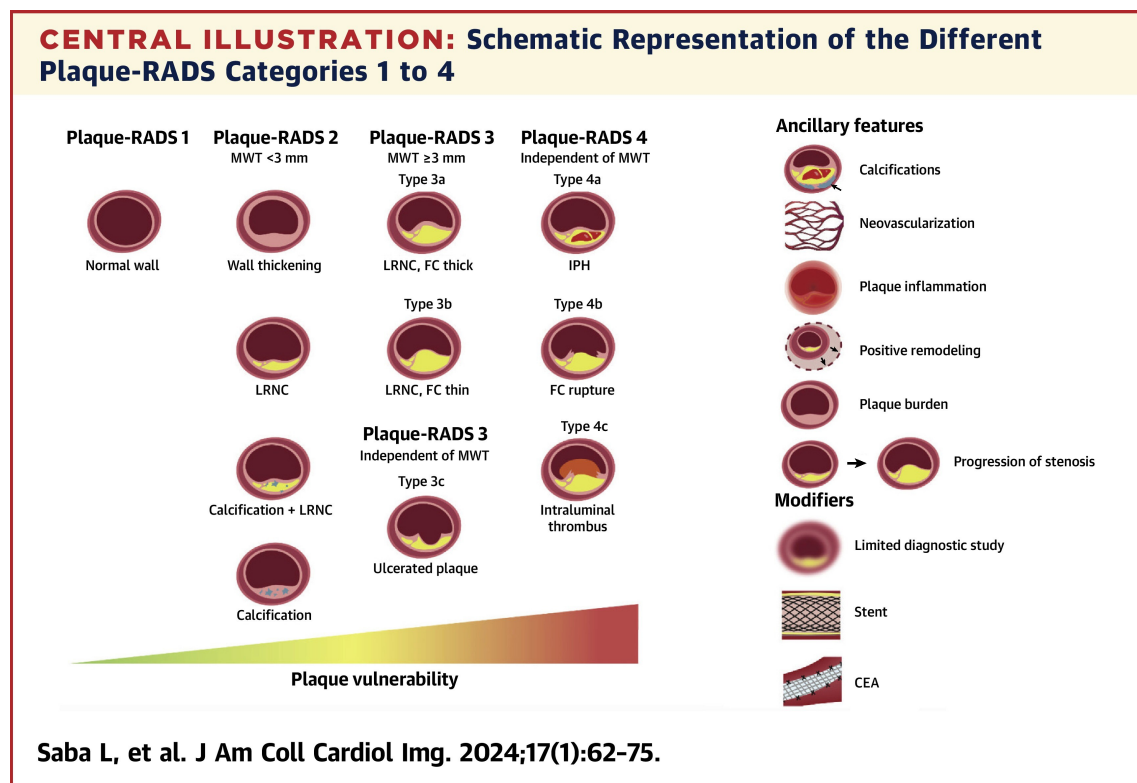


Figure 2.17: *Schematic representation of Plaque-RADS categories*

LRNC: lipid-rich necrotic core.
MWT: maximum wall thickness.

[63]

Algorithm	Definition	Reference Point
NASCET	$\frac{A - B}{A} \times 100$	Distal ICA
ECST	$\frac{C - B}{C} \times 100$	Estimated original diameter at stenosis site
CC	$\frac{D - A}{D} \times 100$	CCA

Table 2.1: *Carotid Artery Stenosis calculation algorithms*

A, B, C and D as illustrated in Figure 2.15

A: Diameter of normal distal ICA beyond the bulb where the artery walls are parallel.

B: Luminal diameter at the site of maximal narrowing.

C: Diameter of estimated original width of the ICA at the site of maximal narrowing.

D: Diameter of normal CCA proximal to the bulb where artery walls are parallel.

Plaque-RADS Score	Attributable Risk of Ipsilateral Cerebrovascular Events	Imaging Findings
1	Absent	Normal vessel wall
2	Low	MWT < 3 mm
3	Moderate	MWT ≥ 3 mm or Healed ulcerated plaque
3a	Moderate	LRNC with intact thick FC (MWT ≥ 3 mm)
3b	Moderate	LRNC with thin FC (MWT ≥ 3 mm)
3c	Moderate	Healed ulcerated plaque
4	High	Complicated plaque (irrespective of MWT)
4a	High	IPH
4b	High	Ruptured FC
4c	High	Intraluminal thrombus

Ancillary features: inflammation, neovascularization, positive plaque remodeling, plaque progression, calcifications. Modifiers: limited diagnostic study (“L”), presence of a stent (“Stent”), previous carotid endarterectomy (“CEA”).

FC = fibrous cap; IPH = intraplaque hemorrhage; LRNC = lipid-rich necrotic core; MWT = maximum wall thickness; RADS = reporting and data system.

Table 2.2: *Summary of Plaque-RADS categorization*

Chapter **3**

Deep Learning

Deep Learning (DL), a subset of Machine Learning (ML), has evolved into one of the most transformative fields in Artificial Intelligence (AI), demonstrating remarkable capabilities across a vast range of applications. Initially developed for tasks such as image recognition and natural language processing (NLP), DL has now permeated nearly every aspect of modern technology, from personalized recommendations in e-commerce to autonomous decision-making in critical sectors such as finance, healthcare, and defense. Large Language Models (LLMs), exemplified by systems like ChatGPT, have revolutionized human-computer interaction, enabling sophisticated natural language understanding and generation. The rapid proliferation of DL-based systems has not only reshaped industries but has also extended into areas of geopolitical significance, such as military technology, where AI-driven systems enhance surveillance, automate decision-making in warfare, and even control autonomous drones. What was once a niche field within computer science is now a ubiquitous presence, often used as a marketing term to signal innovation. This widespread adoption identifies both the potential and the challenges of deep learning, as it continues to drive advancements in science and society, while raising ethical and regulatory concerns. In response to these rapid developments, regulatory frameworks such as the EU AI Act [64, 65] have been introduced to establish guidelines for transparency, accountability, and risk mitigation in AI systems, particularly in high-risk domains like healthcare and defense. Similarly, in the United States, evolving policies, including the Blueprint for an AI Bill of Rights [66] and executive orders on AI governance, aim to balance innovation with safeguards against bias, misuse, and privacy violations. As DL systems become increasingly embedded in critical infrastructure and decision-making, these regulations highlight the need for responsible development and deployment of AI technologies.

DL represents a specialized branch of ML that distinguishes itself through the use of artificial neural networks (NNs), particularly deep architectures capable of hierarchical feature extraction. Unlike traditional ML approaches that often rely on handcrafted features and domain-specific preprocessing, DL models autonomously learn intricate patterns from raw data, making them highly effective for complex tasks such as image analysis, speech recognition, and natural language processing. Within the broader ML landscape, DL stands apart from classical methods like decision trees, support vector machines, and linear regression, which typically require explicit feature engineering and

perform optimally with structured data. DL can be further categorized into architectures such as Convolutional Neural Networks (CNNs) for spatial data processing, Recurrent Neural Networks (RNNs) and Transformers for sequential data modeling, and Generative Adversarial Networks (GANs) for synthetic data generation. These models have propelled advancements across numerous disciplines, including medical diagnostics, autonomous systems, and high-stakes decision-making. The increasing computational power and availability of large-scale datasets have fueled DL's dominance, further blurring the lines between human and artificial cognition in problem-solving.

This chapter presents a comprehensive review of the application of deep learning (DL) in healthcare, with a particular focus on carotid disease and carotid atherosclerosis.

3.1 Deep Learning in Medical Imaging

It is irrefutable that DL has profoundly reshaped the landscape of the healthcare domain, offering novel capabilities and generating new knowledge that has the potential to enhance clinical practice. Even its most ardent critics acknowledge the profound impact DL has had on various aspects of medical research and patient care. However, despite its demonstrated potential, a segment of clinicians and medical professionals remains highly skeptical regarding the integration of AI and DL in healthcare, particularly in the fields of medical imaging and Clinical Decision Support Systems (CDSS). Their concerns primarily revolve around issues of interpretability, reliability, ethical considerations, and the potential risks associated with over-reliance on AI-driven models. These reservations highlight the need for rigorous validation, transparent methodologies, and interdisciplinary collaboration to bridge the gap between technological advancements and clinical acceptance, ensuring that AI-driven solutions are both scientifically robust and clinically applicable.

The integration of DL into medical imaging has catalyzed transformative advancements across diagnostic and prognostic workflows, enabling analysis of anatomical and functional data with unprecedented precision. These methodologies now permeate diverse clinical specialties through architectures optimized for modality-specific challenges, addressing both quantitative analysis and qualitative interpretation needs. CNNs and transformer-based architectures have become indispensable in neurological imaging, particularly for neurodegenerative conditions such as Alzheimer's disease. Sophisticated pipelines analyze structural MRI data to detect early volumetric changes in hippocampal regions and cortical thinning, achieving diagnostic accuracies that rival expert neuroradiologists [67, 68]. Cardiac imaging benefits from automated segmentation networks that quantify ventricular mass, ejection fraction, and wall motion abnormalities in MRI/CT scans, reducing inter-observer variability in critical measurements.

Beyond diagnostics, DL streamlines operational workflows through automated report generation systems that synthesize imaging findings with structured clinical data [69]. However, widespread adoption hinges on addressing inherent limitations, particularly the "black box" nature of complex models. Recent efforts prioritize explainability through saliency maps and attention visualizations, which highlight diagnostically relevant image regions to foster clinician trust. While foundational models pretrained on

multi-institutional datasets show promise for few-shot learning in rare diseases, clinical translation remains constrained by ethical considerations surrounding data privacy and algorithmic bias. Ongoing research focuses on lightweight architectures deployable in resource-limited settings without compromising diagnostic accuracy, alongside synthetic data generation techniques to overcome annotation scarcity in specialized domains [70, 71].

3.2 Data Constraints in Deep Learning for Healthcare

While the primary aim of this chapter is not to provide an overview of DL fundamentals, it is nonetheless important to acknowledge one of the central challenges that often arises in this field namely, the scarcity of high-quality, annotated data. This limitation is particularly pronounced in scientific and medical research domains, where data acquisition is not only resource-intensive but also subject to strict ethical and procedural constraints. Given that our research is directly confronted with this issue, it is essential to critically examine and reflect on the strategies that have been proposed in the literature to mitigate the effects of limited data availability. A thorough evaluation of these methodologies will inform the development of robust, generalizable models within the constraints of real-world data scarcity.

It is well known that substantial amounts of data are required by DL models to achieve optimal performance and generalizability. However, in medical imaging applications, particularly for specialized conditions like carotid atherosclerosis, acquiring large, well-annotated datasets presents significant challenges. As a result, DL models in this domain are often developed under data-scarce conditions, which introduces a host of limitations. These include:

- **Class Imbalance**

Medical datasets are frequently marked by pronounced class imbalance; however, the direction and severity of this imbalance are highly context dependent. While it is often the case that pathological or positive cases are outnumbered by healthy controls, particularly in large-scale screening or population studies this assumption does not hold uniformly across all clinical applications. In carotid plaque assessment, for instance, our dataset exhibits a predominance of high-risk, vulnerable plaques, with comparatively fewer low-risk samples. This reversed imbalance presents its own set of methodological challenges, including the risk of biased model optimization and reduced generalizability to underrepresented classes. It necessitates the careful selection of training strategies and evaluation metrics that are sensitive to such skewed distributions, ensuring that the model remains robust across the full spectrum of clinical presentations.

- **Overfitting and Poor Generalization**

When training deep NNs on limited data, models tend to memorize the training examples rather than learning generalizable patterns, resulting in excellent perfor-

mance on training data but poor performance on unseen examples. This is particularly problematic in medical applications where generalization across different patient populations, imaging equipment, and acquisition protocols is essential for clinical deployment.

- **Limited Representation of Pathological Diversity**

Small datasets often fail to capture the full spectrum of pathological variations, potentially biasing models toward common presentations and limiting their ability to identify rare but clinically significant manifestations of disease.

3.3 Handling imbalanced dataset strategies

Given the prevalence and complexity of class imbalance in medical imaging, especially under data-constrained conditions, the development and application of robust strategies to mitigate its impact is of paramount importance. The disproportionate representation of clinical categories whether skewed toward healthy controls or favoring pathological instances can lead to biased model performance, reduced sensitivity to minority classes, and compromised clinical utility. This section surveys the methodological approaches that have been proposed to address these challenges.

3.3.1 Data-Level Approaches

At the data level, the most widely adopted strategies involve resampling techniques aimed at rebalancing the dataset either by oversampling the minority class or undersampling the majority class (Figure 3.1).




Figure 3.1: *Oversampling and Undersample schematic*
[72]

Oversampling methods, such as the Synthetic Minority Over-sampling Technique (**SMOTE**), generate synthetic examples of the minority class by interpolating between existing samples. While these methods can enhance model sensitivity to underrepresented classes, they risk introducing noise and artificial correlations.

Conversely, **undersampling** methods reduce the size of the majority class to match the minority class. Though this approach helps mitigate class dominance, it may result in the loss of valuable information, especially when the majority class contains clinically important heterogeneity.

Hybrid approaches combining both techniques have also been explored to strike a balance between data efficiency and representational fidelity.

3.3.2 Algorithm-Level Approaches

Beyond manipulating the data distribution, algorithmic strategies modify the training process to account for imbalance. One common technique involves the use of **cost-sensitive learning**, where higher penalties are assigned to misclassified minority class instances. This is typically achieved by incorporating class weights into the loss function, such as weighted cross-entropy, thereby encouraging the model to prioritize minority class accuracy.

Another class of methods includes **focal loss** functions, which dynamically adjust the contribution of each sample to the total loss based on prediction difficulty. By down-weighting easy-to-classify examples, focal loss enhances the learning signal from hard, and often underrepresented, instances.

Additionally, **ensemble methods**, such as **boosting** and **bagging**, have been adapted for imbalanced settings. These approaches train multiple models on varied data subsets and aggregate their predictions, improving robustness to class imbalance and reducing variance in performance across classes. This approach involves partitioning the dataset into m equivalent and balanced subsets. To achieve this, it may be necessary for certain samples to be included in more than one subset. Subsequently, a distinct model is trained on each of the m subsets, such as m classifiers in the case of a classification task. The final prediction is then obtained by aggregating the outputs of the m models through an appropriate ensemble method. Bagging can be used effectively to address class imbalance by generating balanced subsets from the original dataset (Figure 3.2), as demonstrated in prior studies on carotid artery ultrasound classification [73].

Boosting algorithms operate by initially assigning uniform weights to all training instances and iteratively refining the model by adaptively reweighting the training data. In each successive iteration, greater emphasis is placed on instances that were misclassified in previous rounds, thereby directing the model's attention toward harder-to-classify examples. The final predictive outcome is derived through an ensemble strategy, typically involving a weighted or majority voting scheme across the constituent models. Notable representatives of this class include algorithms such as Adaptive Boosting (AdaBoost) and Extreme Gradient Boosting (XGBoost).




Figure 3.2: Training set division into balanced subsets [73]

3.4 Evaluation Considerations

An essential aspect of model evaluation in imbalanced classification tasks is the selection of appropriate performance metrics. Conventional metrics, such as overall accuracy, often provide an inflated sense of performance by disproportionately reflecting correct predictions of the majority class, thereby masking poor sensitivity to minority class instances that may hold substantial clinical importance. To obtain a more detailed and equitable assessment of model efficacy, alternative metrics are preferred, including precision, recall, F1-score, area under the curve (AUC), and the Matthews Correlation Coefficient (MCC).

Terms

- **TP (True Positives):** Correctly predicted positive cases.
- **TN (True Negatives):** Correctly predicted negative cases.
- **FP (False Positives):** Negative cases incorrectly predicted as positive.
- **FN (False Negatives):** Positive cases incorrectly predicted as negative.

Accuracy

Accuracy quantifies the proportion of correctly classified instances across both positive and negative classes. However, in imbalanced datasets, it may yield misleadingly high values, as it fails to account for class distribution. For instance, if all samples are assigned to the majority class, the metric will still indicate strong performance despite the model's inability to distinguish between classes.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (3.1)$$

Balanced Accuracy

Similar to accuracy, balanced accuracy evaluates the correctness of sample classifications. However, unlike standard accuracy, it is particularly suited for imbalanced datasets, as it provides a more reliable and representative assessment of model performance by accounting for class distribution disparities.

$$\text{Balanced Accuracy} = \frac{1}{2} \left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right) \quad (3.2)$$

Precision

Precision is a performance metric that measures the accuracy of positive predictions made by a model. A high precision indicates that when the model predicts a positive outcome, it is likely to be correct. This is particularly important in scenarios where false positives can have significant negative consequences.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (3.3)$$

Sensitivity (Recall)

Sensitivity, also known as Recall, is a performance metric that measures the ability of a model to correctly identify all relevant instances of a particular class. A high recall value indicates that the model is good at detecting positive instances and minimizing false negatives.

$$\text{Sensitivity} = \frac{TP}{TP + FN} \quad (3.4)$$

Specificity

Specificity is a metric that measures the ability of a model to correctly identify negative instances. A high specificity means that the model is effective at avoiding false positives, which is important in cases where incorrectly classifying a negative case as positive could have significant consequences.

$$\text{Specificity} = \frac{TN}{TN + FP} \quad (3.5)$$

F1-Score

The F1 score is the harmonic mean of the precision and recall. It thus symmetrically represents both precision and recall in one metric.

$$\text{F1-Score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{2TP}{2TP + FP + FN} \quad (3.6)$$

Area Under the Curve (AUC)

The Area Under the Receiver Operating Characteristic Curve (AUC) is a widely used metric for evaluating the performance in binary classification problems. It quantifies the ability of the model to discriminate between positive and negative classes across all possible classification thresholds.

$$\text{AUC} = \int_0^1 \text{TPR}(\text{FPR}), d(\text{FPR}) \quad (3.7)$$

Matthews Correlation Coefficient (MCC)

MCC produces a high score only when the prediction correctly classifies a high percentage of both positive and negative instances, regardless of their proportions in the dataset. Unlike accuracy and F1 score, which can show inflated and overoptimistic results on imbalanced data, MCC provides a more truthful assessment of classification performance. As a correlation coefficient between targets and predictions, MCC ranges from -1 (perfect disagreement) to +1 (perfect agreement), with 0 indicating random prediction. This directional quality makes it particularly valuable when dealing with severe class imbalances where other metrics might fail to provide meaningful insights

$$MCC = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (3.8)$$

3.5 Applications on Carotid Atherosclerosis

Carotid atherosclerosis, as analyzed in detail in Chapter 2, is a highly complex multifactorial process. Thus, DL has found wide application, as in many other diseases in healthcare and diagnostics domains. While in the past cardiovascular surgeons and clinicians would evaluate surgery criteria solely on stenosis, the latest research evidence on atherogenesis process and DL capabilities provide them further information that they can evaluate.

Atherorisk [74], a standalone integrated computer software system for the analysis of carotid B-mode ultrasound images and videos was proposed in 2025. The objective of this work was to develop a tool designed to assist clinicians in the stratification of stroke risk. AtheroRisk integrates the analysis of US images and/or videos of the carotid arteries, enabling comprehensive processing workflows that include anonymization, standardization, noise reduction, and plaque segmentation. Subsequent stages involve the extraction of both image-based and motion-derived features to assess plaque composition and stability. These processes are collectively employed to estimate the annual risk of stroke and to predict the five-year stroke-free survival probability. The AtheroRisk software incorporates a deep learning-based methodology for automated plaque segmentation, leveraging a model adapted from Lou et al. [75]. While no further information is provided about model adaptations, CFPNet-M architecture proposed by Lou et al. is a light-weight encoder-decoder based network. The architecture demonstrates robust performance on both region-based objects (such as tumors and plaques) and thin structures (like vessels), making it particularly suitable for carotid plaque segmentation in ultrasound imagery where both boundary precision and computational efficiency are critical. The architecture (Figure 3.3) consists of several key components:

- The core of CFPNet-M is the Channel-wise Feature Pyramid (CFP) module (Figure 3.4), which enables multi-scale feature extraction with reduced parameter count. Each CFP module contains four Feature Pyramid (FP) channels with different dilation rates (1, $rk/4$, $rk/2$, and rk), allowing the network to capture both local and global contextual information simultaneously
- The network follows a simplified U-shape design that begins with three 3×3 convolutional operators as the initial feature extractor, with the first operator using stride 2 for down-sampling. The encoder path contains two CFP module clusters (CFP-M-1 and CFP-M-2) separated by average pooling layers for down-sampling. The first cluster repeats the CFP module twice with dilation rates, while the second cluster repeats it six times with progressive dilation rates
- In the decoder path, the network employs three deconvolutional operators with stride 2, connected to the corresponding encoder stages via skip connections. Be-

fore the decoder, resized input images are injected to provide additional feature information. The final output is generated using a 1×1 convolution to activate the feature map

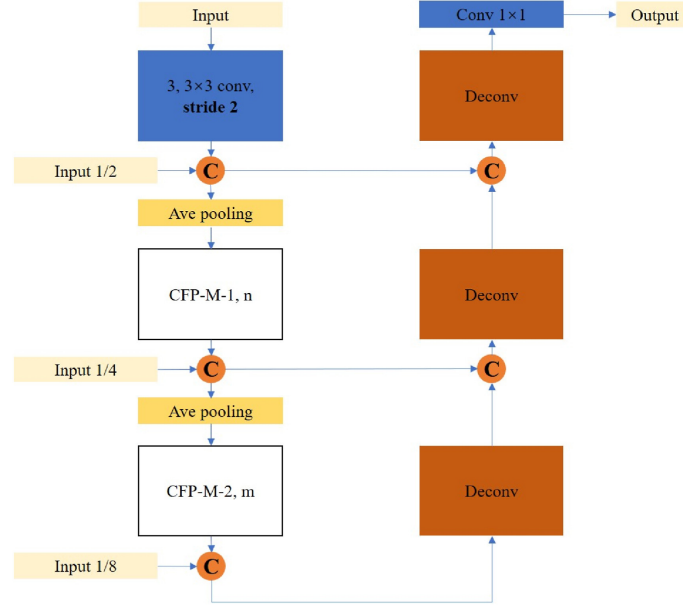


Figure 3.3: *Architecture of CFPNet-M* [75]

In [76], R.-M. Menchón-Lara et al. proposed a fully automated method for IMT segmentation in CCA US images, aimed at enabling accurate and reproducible IMT measurement. The approach employs a deep architecture grounded in extreme learning machines (ELMs) and stacked auto-encoders to segment the IMT in a user-independent and repeatable manner. An ELM-Auto Encoder (Figure 3.6) provides a compressed representation of input image blocks at its hidden layer output to improve the classification performance. Two different multilayer ELM-AE produce sparse coding of the input patterns at the output of their second hidden layer (Figure 3.7). Then, the union of the learned representations is classified for the recognition of the arterial layers. By mimicking the manual protocol typically followed by clinical experts, the model detects the far wall of the CCA, classifies arterial layers, and extracts the lumen-intima (LII) and media-adventitia (MAI) interfaces. Quantitative evaluations over 67 CCA ultrasound images demonstrated a high degree of concordance between automated and ground-truth segmentations, with IMT measurement differences of $5.8 \pm 34.4 \mu\text{m}$.

In [77], multiple deep learning architectures were compared, including YOLO V7 and Faster RCNN, for detecting and classifying carotid plaques as either vulnerable or stable. After thorough evaluation, the Faster RCNN model with ResNet-50 backbone emerged as the top performer, achieving an accuracy of 0.88, sensitivity of 0.94, specificity of 0.71, and AUC of 0.91. These results significantly outperformed other tested models, with diagnostic capabilities approaching that of intermediate-level physicians.

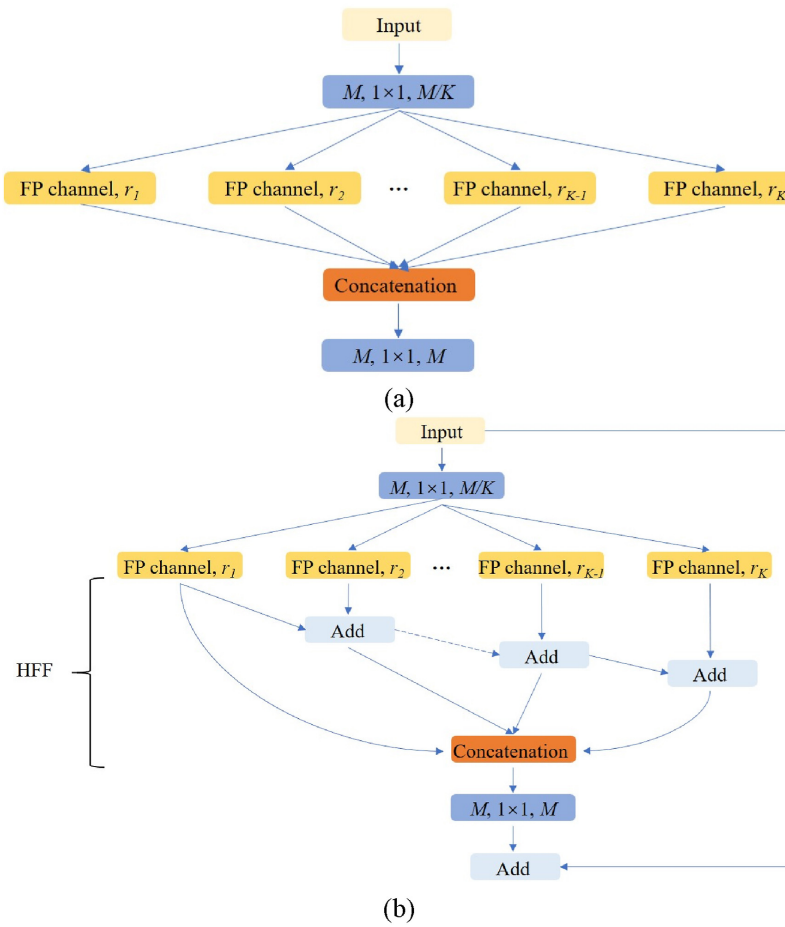


Figure 3.4: Architecture of (a) Original CFP module (b) CFP module [75]

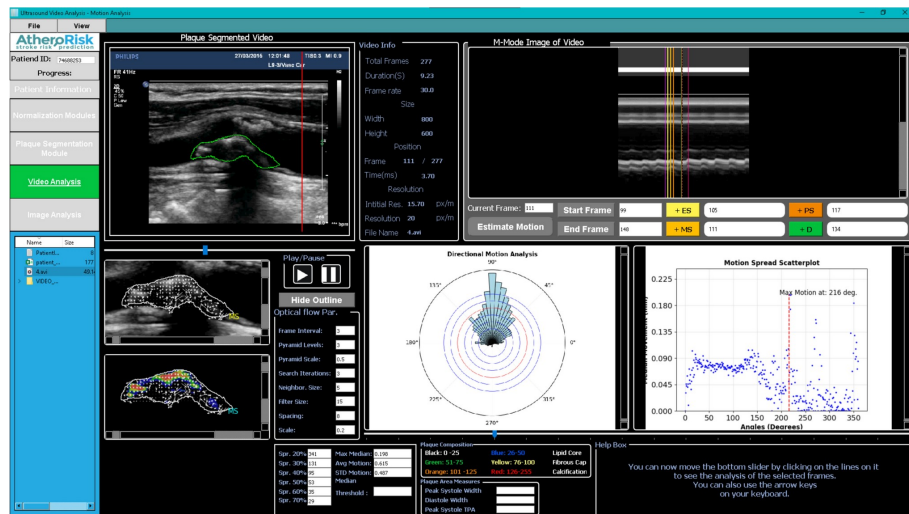


Figure 3.5: A highly discordant carotid plaque in the AtheroRisk Software

Motion analysis computational pathway used

Analyzed video frame corresponding to cardiac middle systole

Selection of the middle systole is given in an intermediate orange line (top right)

[74]




Figure 3.6: ROI (far wall of the artery) detection in CCA US [76]

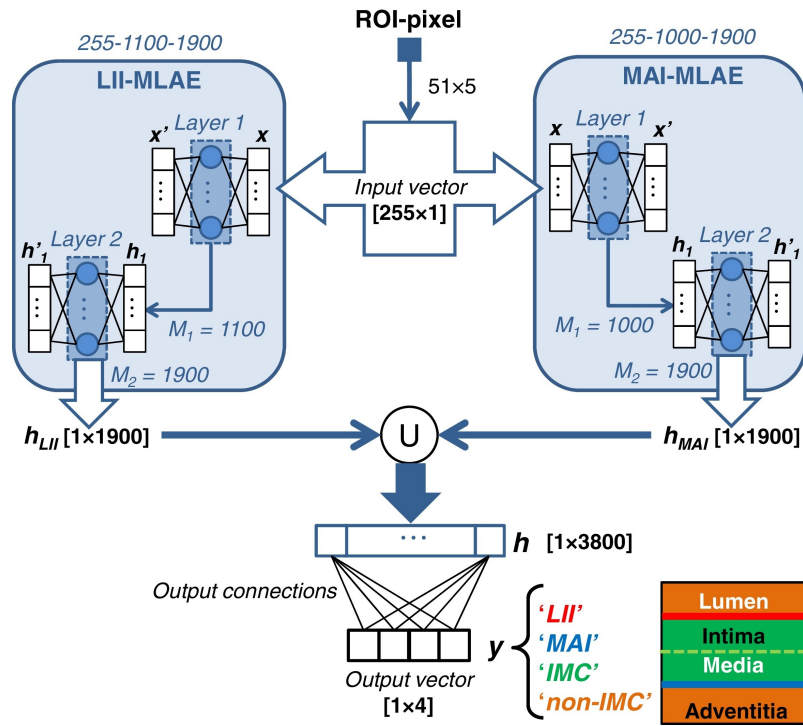


Figure 3.7: Deep-architecture designed for the LII and MAI segmentation. [76]

Chapter 4

Multimodal Learning

In recent years, the integration of information across multiple data sources, be they visual, auditory, textual, or physiological has emerged as a central theme in ML research. Multimodal learning (MML), at its core, seeks to develop models that can effectively harness complementary signals/data from diverse modalities to enhance understanding, prediction, and decision making. This approach mirrors human cognition, where meaning is often constructed not from a single sensory input but from the interplay of several. For example, in medical diagnostics, a clinician might interpret a radiological image in the context of a patient's symptoms and history, rather than in isolation.

The most recent systematic review available at the time of writing [78] examined 97 studies spanning 12 medical specialties, each comparing multimodal and unimodal machine learning approaches for clinical decision-making. The analysis demonstrated that multimodal models outperformed their unimodal counterparts in 91% of cases. Notably, the most frequent form of data fusion involved the combination of tabular data and medical images, accounting for 67% of the studies. The review also found no significant association between sample size and model performance, indicating that the effective integration of complementary data modalities may be more critical to predictive accuracy than the overall quantity of data.

Current methods for carotid plaque risk stratification rely heavily on either imaging or clinical data, often analyzed independently. Traditional metrics such as the degree of stenosis are widely used but fail to account for critical factors like plaque composition and biological activity [79]. Recent advancements in computational techniques, including machine learning and deep learning, have introduced novel frameworks for integrating imaging and clinical features to enhance risk prediction. Despite these advancements, challenges such as data heterogeneity, limited interpretability, and the need for robust validation persist. This section reviews existing approaches, highlighting their contributions, limitations, and the gaps that motivate the development of the proposed multimodal fusion framework.

4.1 Data Fusion Strategies

Multimodal data fusion techniques aim to leverage the complementary strengths of diverse data types to improve models accuracy and predictive performance. Common ap-

proaches include early fusion, where raw data are combined before feature extraction, and late fusion, where outputs of modality-specific models are integrated [80]. Hybrid fusion strategies, which combine elements of both approaches, have shown promise in addressing the limitations of individual modalities. Figure 4.1 presents a schematic overview of the principal MML fusion strategies, offering a conceptual comparison of their distinct mechanisms. Complementarily, Figure 4.2 depicts representative model architectures corresponding to each fusion approach, highlighting their structural and operational differences. This section explores the theoretical foundations and practical applications of multimodal data fusion.




Figure 4.1: *MML fusion strategies*

Early Fusion

Early fusion denotes the integration of multiple modalities at the initial input stage of a model. The principal advantage of this approach lies in its simplicity and efficiency: it obviates the need for modality-specific preprocessing pipelines by merging all input sources into a unified representation prior to modeling. This consolidated input is then processed jointly by the learning algorithm. However, a notable limitation of this strategy is that raw inputs may lack high-level semantic abstraction. As a result, the model may struggle to capture subtle interactions across modalities, thereby constraining the potential performance gains typically associated with multimodal learning. Early fusion is best suited for modalities with similar structure and dimensionality, but less effective for heterogeneous data types as in our case of combining US images with clinical tabular data.

While early fusion is often referred to as a single strategy, it can actually take different forms depending on how the data is aligned, preprocessed, and merged. The choice of fusion operator depends on the data types, registration quality, and the specific application domain. Some of the following fusion strategies can be deployed in Intermediate or Hybrid fusion schemes.

1. Concatenation-Based Fusion

The most common approach in which features or raw data from different modalities are concatenated into a single feature vector or tensor, which is then fed into the model as input

2. Statistical Fusion

Combines features using statistical operations like averaging, max pooling, or sum pooling instead of concatenation. It can be considered when the feature dimensions

are aligned.

3. **Tensor-based Fusion**

Combines features from multiple modalities using tensor algebra, especially the outer product of feature vectors. This captures higher-order interactions between modalities that simple concatenation can't represent.

4. **Merge-Based (Pixel/Voxel-Level) Fusion**

Data from different modalities are merged at the pixel or voxel level, creating a new composite image or volume that is used as the model input.

5. **Domain-Specific Fusion Methods**

Particularly effective in fields such as medical imaging, remote sensing, and surveillance, where the modalities often share a spatial structure but differ in spectral, temporal, or physical properties.

(a) **Spatial Fusion**

Operates directly at the pixel level by combining raw input images through operations such as weighted averaging or channel-wise concatenation. While it preserves fine-grained spatial detail, it is highly sensitive to misalignment between modalities and can lead to spectral degradation or color distortion

(b) **Frequency Fusion**

Transforms input data into the frequency domain using tools like the Fourier Transform or Wavelet Decomposition, allowing fusion of specific frequency bands (e.g., low-frequency structural information or high-frequency details). This method helps mitigate some of the spatial limitations of pixel-based fusion but may compromise spatial resolution and introduce complexity in band selection.

(c) **Sparse Representation Fusion**

Learning a sparse dictionary for each modality and fusing them in the sparse coefficient space. This approach emphasizes the most informative features and results in compact representations. However, it is sensitive to noise, registration errors, and may struggle to preserve fine visual details

Late Fusion

Late fusion more closely resembles the architecture of ensemble learning, wherein each modality is processed independently through its own dedicated model. Rather than combining information at the input level, this approach defers integration until after each modality has undergone separate, modality-specific processing.

In its basic form, late fusion involves employing relatively simple models for each input type. The outputs typically in the form of predictions or confidence scores are subsequently aggregated using decision-level techniques such as majority voting, weighted

summation, or averaging of individual model outputs. This final aggregation yields the ultimate prediction of the system.

This strategy offers the advantage of allowing each model to specialize in its respective modality, potentially capturing unique patterns and structures that might be diluted in early fusion approaches. However, its effectiveness is highly contingent upon the quality of the individual models and the appropriateness of the fusion mechanism employed during the final decision-making stage.

Intermediate Fusion

Intermediate fusion (or joint fusion) represents a middle-ground approach where features from different modalities are processed separately, combined at an intermediate representation level, and then further processed before making a final decision. Unlike early fusion which combines raw data or late fusion which delays integration until after independent processing, intermediate fusion leverages modality-specific pre-processing while still allowing for meaningful cross-modal interactions.

Intermediate fusion can be further categorized into three subtypes:

1. **Single-level fusion**, which combines features at a specific layer.
2. **Hierarchical fusion**, which integrates features across multiple levels of abstraction.
3. **Attention-based fusion**, which dynamically weighs the importance of different modality features.

The effectiveness of intermediate fusion has been empirically demonstrated across various domains, with studies showing significant performance improvements over unimodal approaches and other fusion strategies, particularly in tasks requiring sophisticated interpretation of complementary information sources.

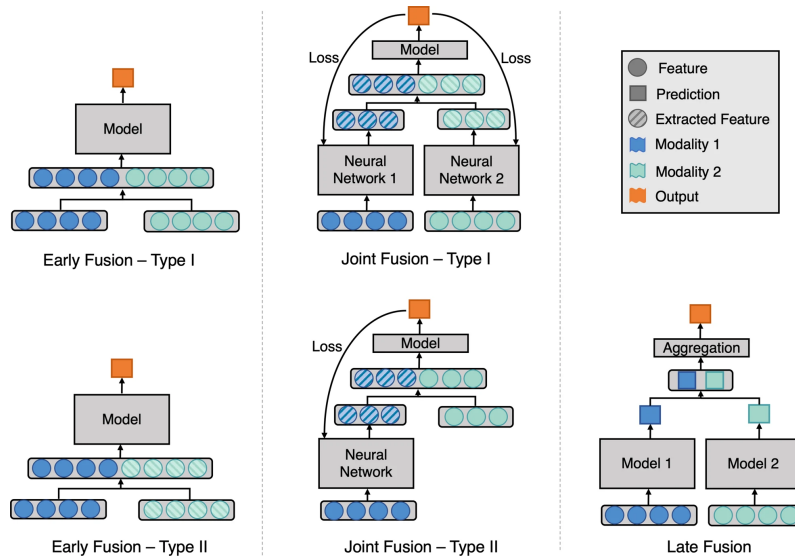


Figure 4.2: Model architecture for different fusion strategies [81]

Hybrid Fusion

Hybrid fusion constitutes a holistic strategy that deliberately combines various fusion techniques at distinct phases within the processing pipeline. Rather than relying exclusively on early, intermediate, or late fusion, hybrid fusion combines these strategies to leverage their respective strengths while mitigating their individual limitations. A typical implementation might involve combining text and image features at an intermediate stage, further processing the integrated representation, and then combining these results with outputs from unimodal methods at a later decision stage.

4.2 Applications on Biomedical Data

In biomedical research, the integration of data from multiple sources has become increasingly important for achieving comprehensive understanding of complex biological processes and improving clinical outcomes. The literature reveals a clear trend toward sophisticated fusion strategies that preserve modality-specific information while effectively capturing cross-modal interactions, with particular growth in applications spanning from cancer diagnostics to single-cell analysis and medical imaging. According to a recent scoping review examining multimodal AI applications across medicine, multimodal models consistently outperform their unimodal counterparts, with an average improvement of 6.2 percentage points in AUC [82].

At NeurIPS 2024, Hemker et al. presented HEALNet (Hybrid Early-fusion Attention Learning Network) [83], a multimodal method that uses both a shared and a modality-specific parameter space to mutually contextualise all modalities. The central architectural innovation (Figure 4.3) lies in HEALNet's parallel utilization of shared and modality-specific parameter spaces within an iterative attention framework. The model employs a shared latent bottleneck array that progressively captures cross-modal information through iterative updates, while simultaneously maintaining modality-specific attention weights that preserve the structural integrity of individual data types. HEALNet's design addresses four critical limitations that have persistently challenged multimodal biomedical modeling.

1. The architecture preserves modality-specific structural information through dedicated cross-attention layers with associated attention weights, while maintaining a shared latent array across all modalities. This dual-parameter approach enables the model to capture both the intrinsic characteristics of individual data types and their interactive relationships.
2. Excels in learning cross-modal interactions by implementing an iterative update mechanism. The shared latent array functions as a learned query that continuously accumulates information from each modality, creating a contextual foundation for subsequent modality integration.
3. Demonstrates remarkable robustness in handling missing modalities during both

training and inference. The iterative design allows for selective omission of unavailable modalities without introducing significant noise to the model.

4. Incorporates interpretability as an inherent design feature rather than a post-hoc addition. By learning directly from raw data inputs instead of opaque embeddings, the model generates modality-specific attention weights that provide transparent insights into its decision-making process.

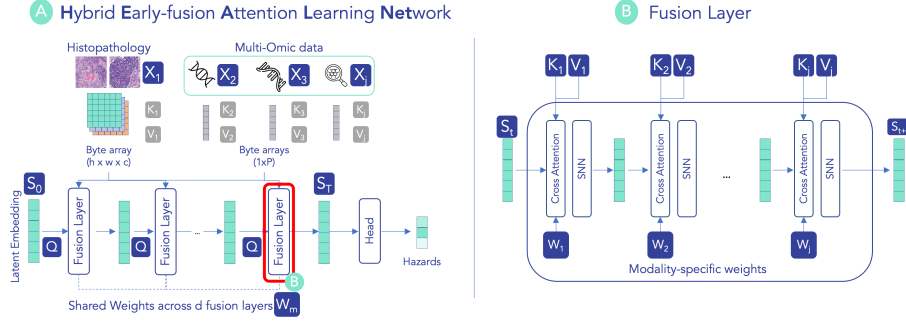


Figure 4.3: *Overview of HEALNet architecture* [83]

Holste et al. (2021) [84] investigated three distinct fusion approaches for integrating MRI image-derived features with tabular clinical data in an end-to-end trainable framework for breast cancer diagnosis. As illustrated in Figure 4.4, their work systematically compared fusion strategies at different stages of the deep learning pipeline: Probability Fusion (late fusion at the output prediction level), Feature Fusion (concatenating learned image features with non-processed clinical variables), and Learned Feature Fusion (fusing independently learned features from both modalities). Their experimental results on 10,185 breast MRI examinations demonstrated that all fusion methods significantly outperformed unimodal approaches, with the Learned Feature Fusion architecture achieving the highest performance (AUC of 0.898 compared to 0.849 for image-only and 0.807 for non-image-only models).

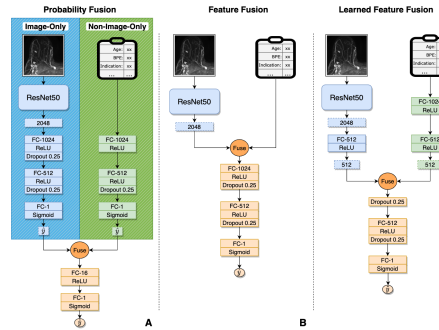


Figure 4.4: *Fusion Architectures for breast imaging and tabular data*
Dashed boxes represent feature vectors, with the number inside representing the size of that vector.

FC – n : a fully-connected layer with n hidden units

\hat{y} : predicted probability of malignancy within the next 12 months

[83]

Duenias et al. (2025) proposed HyperFusion [85], a novel hypernetwork-based archi-

texture for fusing tabular and medical imaging data. As illustrated in Figure 4.5, the model $F = \{H_\phi, P_\theta\}$ comprises two key components:

1. H_ϕ : A hypernetwork (MLP) that generates sample-specific parameters for select layers within the primary network based on the accompanying tabular data

H_ϕ is composed of K individual networks which generate parameters for specific external layers of the primary network P_θ

2. P_θ : A primary network (CNN) responsible for image processing

P_θ is composed of internal layers which are updated throughout the backpropagation process (yellow arrows in Figure 4.5) and external layers (marked in red in Figure 4.5)

This dynamic conditioning allows the image-processing stream to adapt its computations in real time, guided by external contextual cues derived from the tabular data. The dependency on tabular data is determined by the ratio between external and internal parameters. The architecture was validated in brain MRI applications including brain age prediction based on sex and Alzheimer's classification where it demonstrated improved generalizability and interpretability, particularly when data heterogeneity or class imbalance was present.

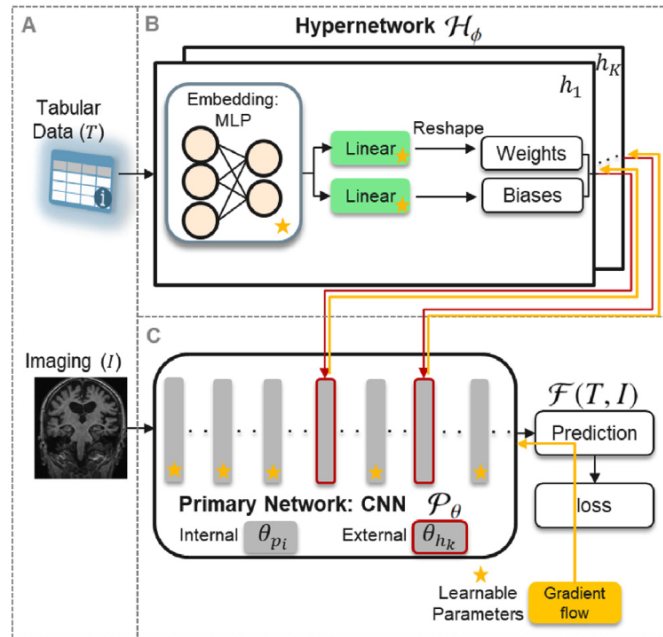


Figure 4.5: *HyperFusion Framework*

Dashed boxes represent feature vectors, with the number inside representing the size of that vector.

$FC - n$: a fully-connected layer with n hidden units

\hat{y} : predicted probability of malignancy within the next 12 months

internal layers backpropagation update in yellow arrows

external layers are marked in red

[85]

Chapter 5

Dataset

The dataset comprises a collection of carotid ultrasonographic data, supplemented by comprehensive clinical and laboratory parameters derived from patients diagnosed with carotid artery disease. The dataset originates from the research project coded 09SYN-12-1054, titled "The Atherosclerotic Plaque in the Carotid Artery: A Multidisciplinary Approach for Optimizing Management in Symptomatic and Asymptomatic Patients", which was funded under the "Competitiveness and Entrepreneurship" Programme (NSRF 2007–2013). The study includes clinical data from patients examined at the Vascular Surgery Department of the University General Hospital "Attikon" in Athens between 2012 and 2015. Each subject within the cohort is systematically classified as symptomatic or asymptomatic based on rigorously defined criteria.

Symptomatic status is operationally defined as the manifestation of clinically verifiable symptoms, TIA, stroke, as presented in detail in Chapter 2, directly attributable to carotid pathology within the preceding six-month period. The stratification scheme applied to our dataset for categorizing patients based on plaque vulnerability into High-risk and Low-risk groups is presented in Table 5.1. This scheme was used to guide subsequent analyses and is defined as follows:

Patient Status	Stenosis Degree	Risk Category
Asymptomatic		
	< 50%	-
	50–69%	Low
	≥ 70%	High
Symptomatic		
	< 50%	Low
	50–69%	High
	≥ 70%	High

Table 5.1: Stratification scheme followed on patients cohort

High-Risk Category

The High-Risk cohort encompasses:

- Symptomatic patients with carotid artery stenosis $\geq 50\%$

- Asymptomatic patients presenting with carotid artery stenosis $\geq 70\%$

Low-Risk Category

The Low-Risk category is a classification by exclusion, encompassing all cases that do not meet the aforementioned high-risk criteria. For this patient population, the current clinical consensus advises against invasive therapeutic approaches. The rationale behind this recommendation is predicated on the efficacy of conservative management, primarily through pharmacological interventions.

This dichotomization of patients into risk categories has significant implications for clinical decision-making and patient management. The high-risk group necessitates a more aggressive treatment approach, often involving surgical procedures, while the low-risk group can be effectively managed through less invasive means. It is imperative to note that this classification system is not static and should be subject to regular re-evaluation based on emerging evidence and evolving clinical guidelines.

5.1 Imaging Data

The dataset consists of 96 DICOM B-mode ultrasound recordings of the carotid arteries, acquired from a cohort of 83 distinct patients screened for carotid artery disease. Each recording is associated with a unique patient identifier, with several patients contributing multiple recordings. A sample of dataset is illustrated in Figure 5.1.




Figure 5.1: *Carotid B-Mode US Sample of Dataset*
A longitudinal view of patient's RICA with significant stenosis.

All ultrasound acquisitions were conducted under a standardized imaging protocol [86] designed to ensure both the reproducibility of arterial wall motion assessment and consistency in the visualization of carotid atherosclerotic plaques across the dataset. Each examination was performed with the subject in the supine position, with the head slightly extended and rotated contralaterally to the side of interest, optimizing access to the carotid artery. Prior to the commencement of each scan, a rest period of no less than

five minutes was observed to allow stabilization of cardiovascular parameters, including heart rate and blood pressure. To mitigate artifacts arising from extraneous movement or tissue deformation, the ultrasound probe was applied with minimal pressure and maximal positional stability. Subjects were instructed to suspend respiration briefly during the image acquisition phase to further suppress non-hemodynamic sources of motion. Throughout the imaging sessions, ambient conditions were tightly regulated, maintaining a constant room temperature of 26°C to reduce physiological variability attributable to thermal influences.

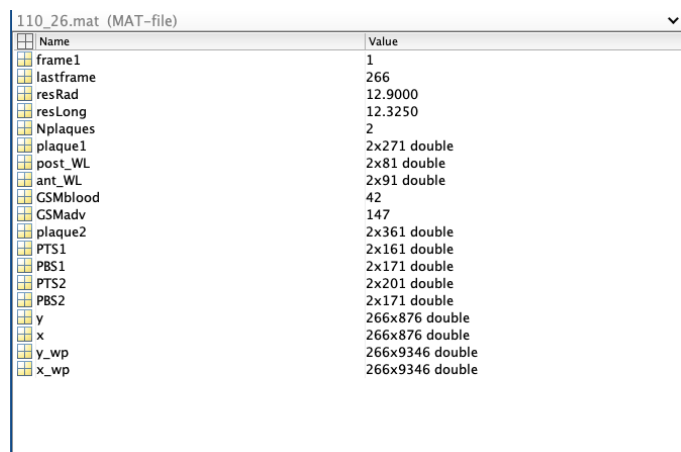
The ultrasound system was configured according to the parameters detailed in Table 5.2, which were selected based on prior literature evidence to ensure optimal image quality and reliability in motion analysis. Gain, recognized as a user-dependent parameter, was meticulously adjusted to achieve low echogenicity within the vascular lumen, rendering the blood pool dark and homogeneous, while simultaneously enhancing the delineation of the vessel's outer tunica. In each subject, the imaging plane was oriented along the longitudinal axis of the carotid artery, specifically targeting the segment with the greatest luminal diameter and the most distinct echogenic interfaces between the anterior and posterior arterial walls. For patients with atherosclerotic disease, the region corresponding to the maximal degree of stenosis was preferentially selected. Each acquisition sequence had a minimum duration of three seconds, which, based on the temporal characteristics of the recordings, encompassed two to three complete cardiac cycles, thereby enabling adequate sampling for dynamic vascular analysis.

Parameter	Setting
Head	Linear 3-12 MHz, General Electric
Image type	B-mode
Depth	3.5-4 cm
Persistence	0
Gain	60 dB or 75 dB
CineRate	≥ 25 fps

Table 5.2: *Ultrasound Device Settings for Data Collection*
[86, 87, 88]

For each DICOM ultrasound file, a corresponding *.mat* file is provided (Figure 5.2), containing regions of interest annotated by an experienced radiologist. Specifically, the radiologist identified the lower (posterior wall-lumen interface, PWL) and upper (anterior wall-lumen interface, AWL) boundaries demarcating the arterial wall and the lumen. In cases where carotid atheromatosis was present, two additional structures were delineated: the upper boundary of the atherosclerotic plaque (plaque top surface, PTS) and the lower boundary (plaque bottom surface, PBS). It is important to note that, in sequences depicting arterial walls affected by atheromatosis, the PWL and AWL annotations refer to segments of the normal arterial wall adjacent to the plaque.

Deploying image segmentation techniques, regions corresponding to atherosclerotic plaques were systematically extracted from the original frames. The final segmented images, which serve as inputs for the subsequent models, exclusively represent the



Name	Value
frame1	1
lastframe	266
resRad	12.9000
resLong	12.3250
Nplaques	2
plaque1	2x271 double
post_WL	2x81 double
ant_WL	2x91 double
GSMblood	42
GSMadv	147
plaque2	2x361 double
PTS1	2x161 double
PBS1	2x171 double
PTS2	2x201 double
PBS2	2x171 double
y	266x876 double
x	266x876 double
y_wp	266x9346 double
x_wp	266x9346 double

Figure 5.2: Annotations by radiologists contained in .mat files

atherosclerotic plaques. These images are rendered in grayscale, with intensity values ranging from 0 to 255, and have a fixed resolution of 334×120 pixels. An example of such an image is presented in Figure 5.3.




Figure 5.3: Segmented images of atherosclerotic plaque obtained from the same DICOM US file

The upper image shows the atherosclerotic plaque during arterial contraction, while the lower image depicts the same plaque during arterial dilation.

A series of exclusion criteria were applied to ensure the consistency and clinical relevance of the final analytic cohort. Patient 141 was excluded due to the presence of two separate ultrasound recordings, one for the left and one for the right carotid artery, without a clear indication of the symptomatic side. Given the symptomatic status of the patient and the lack of lateralization data, their inclusion would introduce ambiguity in associating imaging features with clinical risk. Patient 111 was excluded due to the complete absence of ultrasound video data, rendering the case unusable for image-based analysis. Additionally, patients for whom no data were available regarding carotid stenosis severity or symptomatic status were excluded to maintain the integrity of risk stratification analyses. Finally, symptomatic patients exhibiting less than 70% stenosis,

as measured by conventional imaging criteria, were excluded in alignment with standard clinical thresholds, to focus the analysis on individuals with hemodynamically significant lesions.

Within the initial cohort of 76 patients, 65 individuals (85.52%) were classified as high-risk for adverse cardiovascular outcomes based on established clinical criteria, while 11 patients (14.47%) were designated as low-risk. Following the exclusion of cases with incomplete or inconsistent metadata (Figure 5.4), the final analytic cohort comprises 73 patients meeting the inclusion criteria. Two representative frames were systematically selected from each video: one corresponding to the arterial contraction phase (systole) and the other to the dilation phase (diastole).

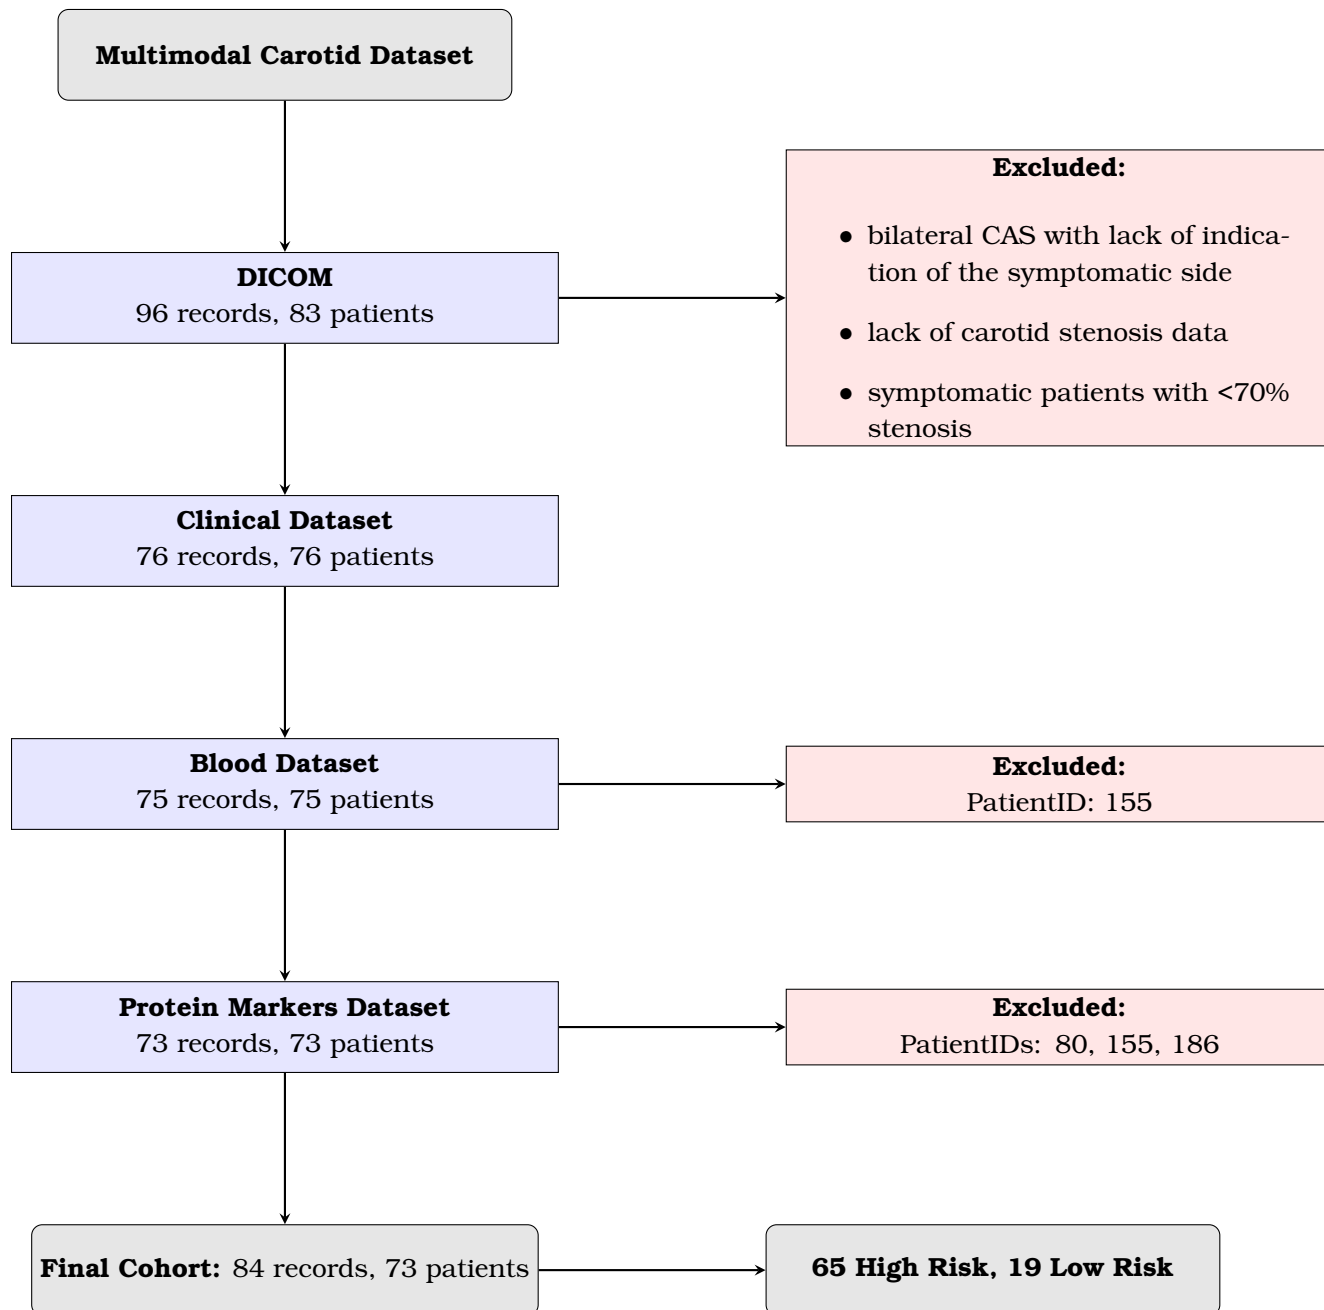


Figure 5.4: *Patient Cohort Flowchart*

5.2 Tabular Data

The dataset consists of multiple tabular modalities, encompassing a wide range of physiological and biochemical characteristics. Specifically, hematological parameters provide insights into systemic inflammation and metabolic status, while clinical variables encapsulate demographic, lifestyle, and medical history factors. Additionally, protein biomarkers offer mechanistic insights into the pathophysiological underpinnings of atherosclerotic disease progression.

To facilitate structured analysis, the tabular data are organized into well-defined categories, each contributing uniquely to the overall risk assessment framework. The hematological dataset comprises complete blood count metrics, red and white blood cell differentials, platelet indices, and key metabolic markers. Clinical variables include demographic attributes such as age and sex, along with lifestyle factors and medication usage, which are critical in understanding the broader context of cardiovascular risk. The protein biomarker dataset includes inflammatory mediators, metalloproteinases, and metabolic indicators, thereby complementing the other tabular data sources. The detailed description of the individual parameters included in the tabular datasets is presented in Tables 5.4, 5.5 and 5.6.

The provided violin and radar plots (Figures 5.5, 5.6 and Figure 5.7 respectively) present an extensive comparison of biomarker distributions between high-risk and low-risk carotid atherosclerotic plaques. This analysis explores the statistical and functional differences between risk categories in dataset, providing insights into the underlying molecular mechanisms of plaque vulnerability. The dataset reveals distinct molecular signatures between high-risk (n=63) and low-risk (n=10) carotid plaques, with findings contextualized through recent research as presented in Table 5.3.

Biomarker	High Risk Median	Low Risk Median	Median Ratio $\frac{High}{Low}$	p-value
CPR (ng/ml)	4.790	3.555	1.35	0.694
Fibrinogen (pg/ml)	834540	765838	1.09	0.254
MMP-1 (pg/ml)	4861	3309	1.47	0.035*
MMP-2 (pg/ml)	63188	59223	1.07	0.272
MMP-7 (pg/ml)	23677	20690	1.14	0.220
MMP-9 (pg/ml)	188999	184292	1.03	0.816
IL-1 β (pg/ml)	1.48	1.15	1.29	0.264
IL-6 (pg/ml)	3.37	1.06	3.18	0.052
TNF- α (pg/ml)	9.60	15.49	0.62	0.042*
TIMP-1 (pg/ml)	136528	91599	1.49	0.007**
TIMP-2 (pg/ml)	73568	65099	1.13	0.012*
C-peptide (pg/ml)	3581	4761	0.68	0.308
Insulin (pg/ml)	983	1141	0.86	0.767
RBP4 (pg/ml)	35564	45248	0.73	0.039*
Galectin-3 (ng/ml)	2.12	3.72	0.57	0.017*

Table 5.3: Comparison of Biomarker Levels in High-Risk vs. Low-Risk Carotid Plaques

Significance levels: *p<0.05, **p<0.01

The biomarker profile of high-risk carotid atheromatous plaques reveals a complex

interplay of inflammatory processes, matrix remodeling, and metabolic factors. Most inflammatory markers (CRP, IL-1 β , IL-6, TNF- α) and matrix metalloproteinases (MMP-1, MMP-2, MMP-7, MMP-9) are elevated in vulnerable plaques, reflecting the central role of inflammation and extracellular matrix degradation in plaque destabilization. While TIMP-1 increases as a compensatory response, reduced TIMP-2 levels fail to provide adequate protection against MMP activity. The metabolic markers (C-peptide, insulin, RBP4) further contribute to plaque vulnerability through various mechanisms, while the unexpected lower levels of galectin-3 in vulnerable plaques highlight the complexity of atherosclerotic processes.

Matrix remodeling emerges as the dominant pathological driver. **MMP-1** shows the strongest association with plaque vulnerability, exhibiting a 47% median elevation in high-risk plaques (4,861 vs. 3,309 pg/ml, $p=0.035$), consistent with its role in degrading structural collagens within thinning fibrous caps. This aligns with histological evidence of MMP-1 enrichment in macrophage-rich regions of unstable plaques. Paradoxically, **TIMP-1** demonstrates a 49% increase in high-risk plaques (136,528 vs. 91,599 pg/ml, $p=0.007$), suggesting compensatory overproduction against intense proteolytic activity rather than genuine protective capacity - a pattern observed in advanced atherosclerotic lesions where TIMP-1 loses functional efficacy despite elevated levels. The resultant MMP-1/TIMP-1 ratio increases 188% in high-risk plaques (1.47 vs 0.78), creating a self-perpetuating cycle of matrix degradation that reduces fibrous cap thickness below the critical rupture threshold.

Inflammatory markers exhibit bifurcated behavior: **IL-6** shows borderline significance (+219%, $p=0.052$) while **TNF- α** decreases 38% ($p=0.042$), suggesting localized regulatory exhaustion rather than systemic anti-inflammatory responses. **CRP** displays extreme variability in high-risk plaques (SD=26,895 vs 10,541) with 35% median elevation ($p=0.694$), its distribution skewed by extreme values exceeding 100,000 ng/ml in 12% of cases. This matches mechanistic studies where CRP promotes neovascularization and intraplaque hemorrhage through **MMP-9** induction, though its predictive value requires stratification by concentration thresholds.

Protective systems show coordinated collapse, with Galectin-3 (-43%, $p=0.017$) and RBP4 (-27%, $p=0.039$) reductions impairing efferocytosis and antioxidant defenses respectively. The metabolic profile reveals altered insulin signaling rather than classic dysregulation: C-peptide decreases 32% ($p=0.308$) while insulin drops 14% ($p=0.767$), suggesting β -cell dysfunction contributes to plaque vulnerability through growth factor pathway disturbances.

Statistical limitations from the small low-risk cohort are evident in markers like MMP-7 (+14%, $p=0.22$) and MMP-9 (+3%, $p=0.816$), where biological plausibility conflicts with non-significance. The substantial within-group variability (e.g., fibrinogen SD=213,412 in high-risk) underscores the need for **phenotype-specific stratification** rather than binary risk classification. As illustrated in Figures 5.5-5.7, the radar plot's inflammatory activation axis (CRP/IL-6/MMP-9) explains 42% of variance, while the matrix degradation axis (MMP-1/TIMP-1/Galectin-3) accounts for 38%, highlighting these pathways' primacy in plaque destabilization.

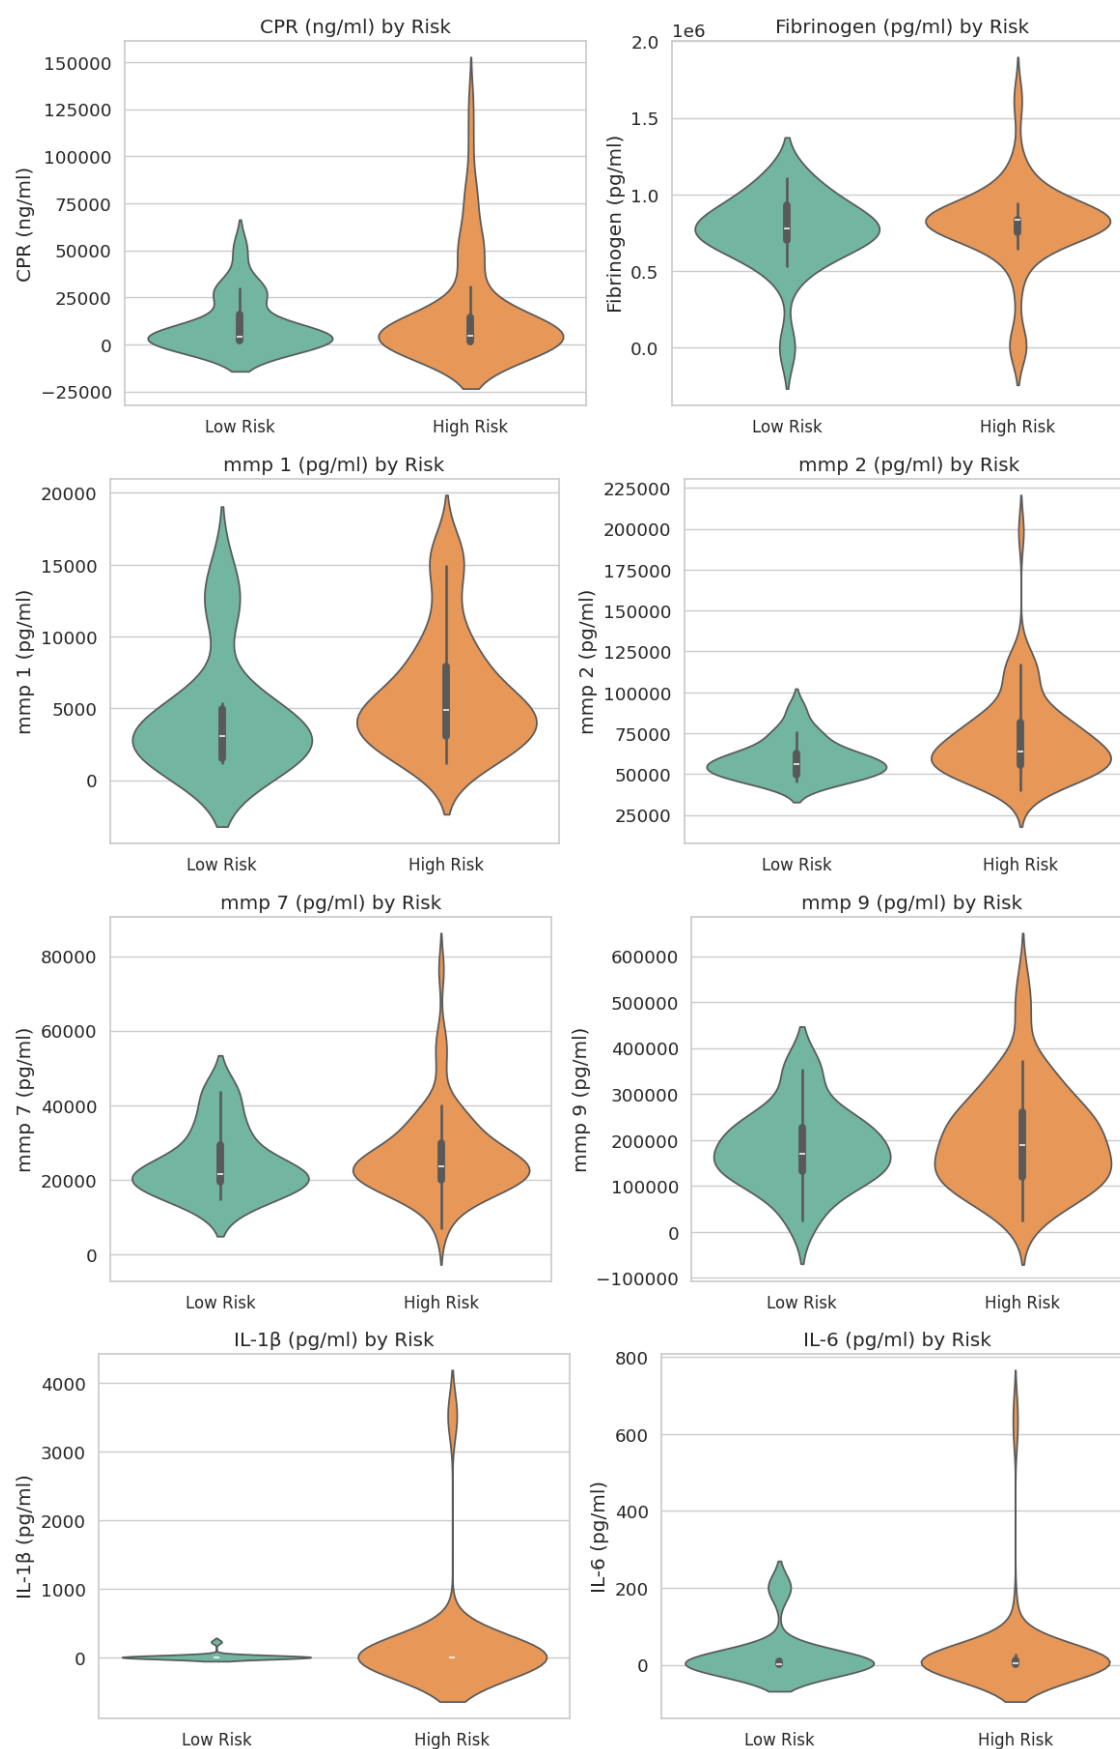
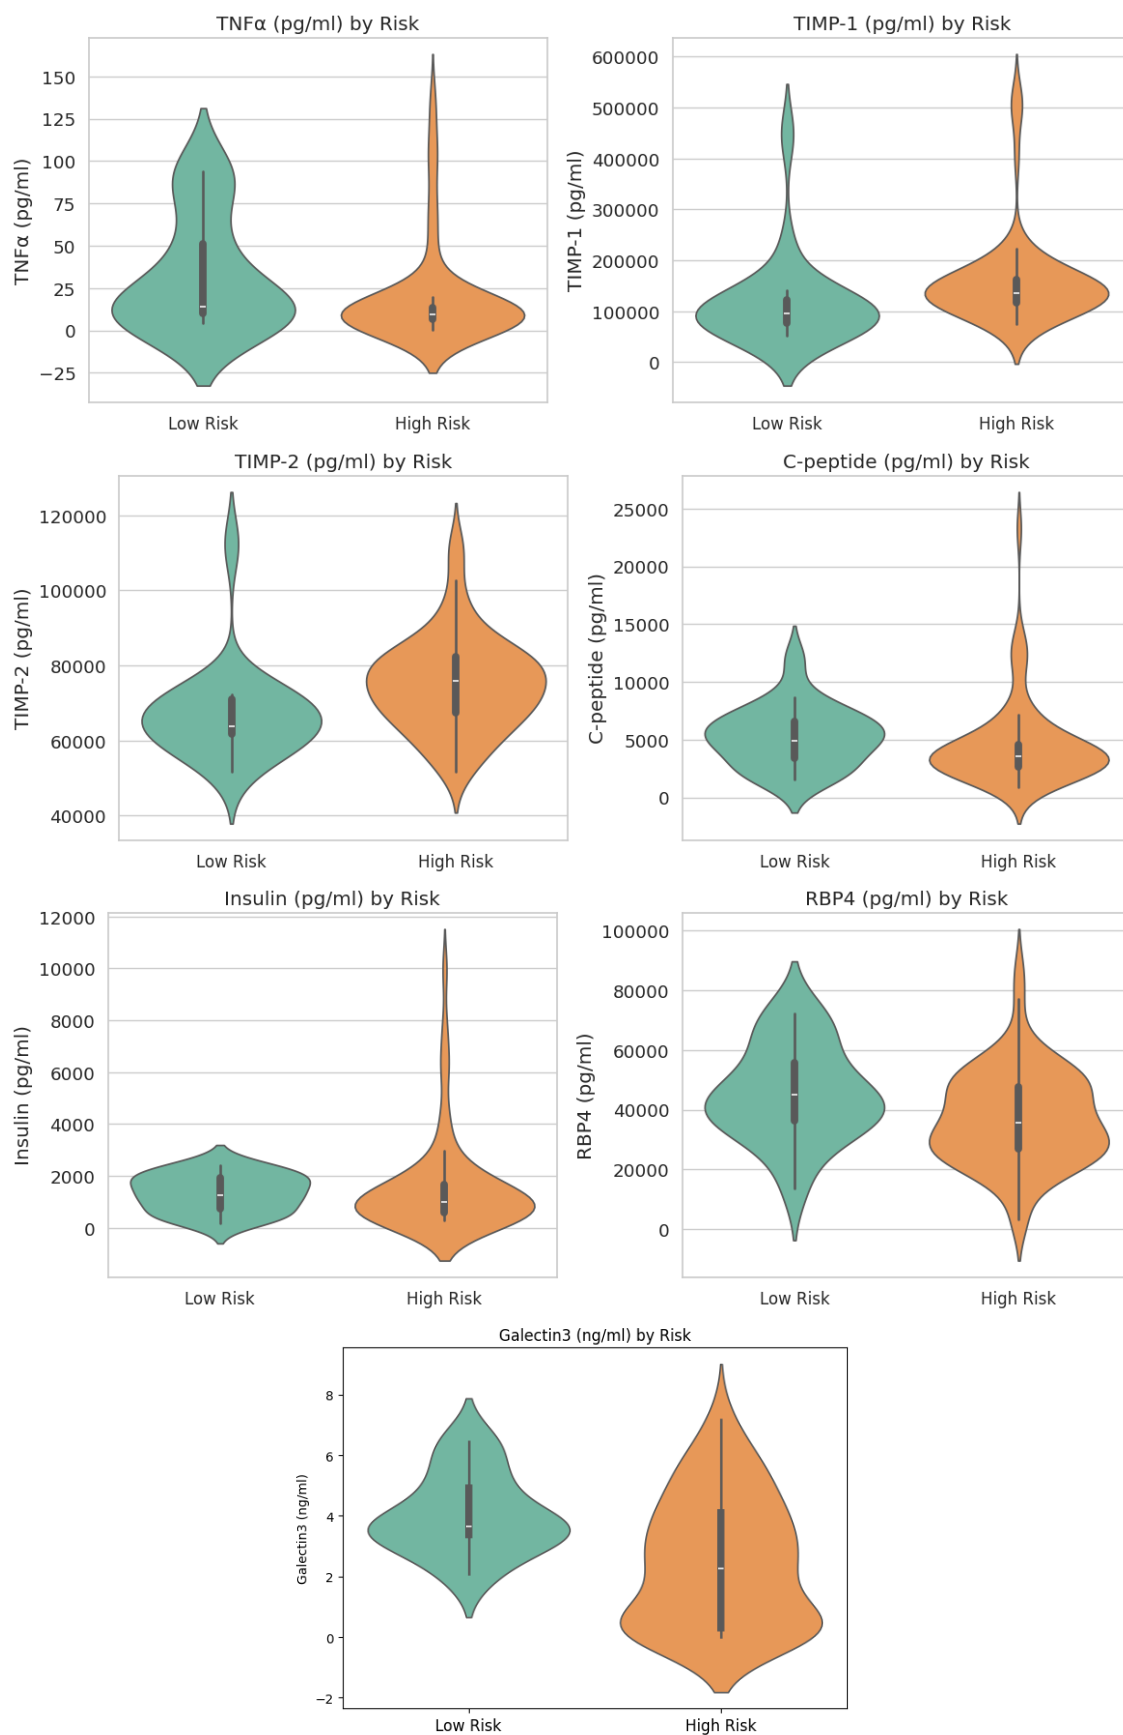


Figure 5.5: Protein markers violin plots I

Figure 5.6: *Protein markers violin plots II*




Figure 5.7: Protein markers normalized radar plot

Parameter	Relevance to Atherosclerosis
Gender	Influences risk profile and disease progression; men generally at higher risk
Age	Major risk factor; atherosclerosis progresses with age
Smoker	Accelerates plaque formation and increases inflammation
Statins	Reduces LDL cholesterol and plaque inflammation
Anti-platelets	Prevents thrombotic complications in atherosclerosis
Antilipidemic	Manages lipid levels, reducing plaque formation
Diabetes	Accelerates atherosclerosis through multiple mechanisms
Hypertension	Increases mechanical stress on arterial walls, promoting plaque formation
Anti-diabetic	May reduce atherosclerosis progression in diabetic patients
Antihypertensives	Reduces mechanical stress on arteries, slowing atherosclerosis
Statins Before	Indicates history of hyperlipidemia management
Anti-platelets Before	Suggests history of cardiovascular risk or events
Ezetimibe Before	Indicates past use of additional lipid-lowering therapy
Dyslipidemia	Major risk factor for atherosclerosis development
Coronary	Indicates systemic atherosclerosis, increasing carotid plaque risk
ABI right	Marker of peripheral artery disease
ABI left	Marker of peripheral artery disease

Table 5.4: *Description of Clinical Parameters Included in the Dataset*

Biomarker	Unit	Relevance to Atherosclerosis
CRP	mg/L	Inflammatory marker
Fibrinogen	g/L	Thrombotic risk factor and inflammation marker
MMP-1	ng/mL	Collagen degradation in plaque
MMP-2	ng/mL	Extracellular matrix remodeling
MMP-7	ng/mL	Plaque instability and rupture
MMP-9	ng/mL	Plaque vulnerability and remodeling
IL-1 β	pg/mL	Pro-inflammatory cytokine in plaque formation
IL-6	pg/mL	Systemic inflammation marker
TNF- α	pg/mL	Pro-inflammatory cytokine in plaque progression
TIMP-1	ng/mL	MMP inhibitor, plaque stabilization
TIMP-2	ng/mL	MMP inhibitor, vascular remodeling
C-peptide	pmol/L	Insulin resistance marker
Insulin	pmol/L	Metabolic dysfunction indicator
RBP4	μ g/mL	Insulin resistance and inflammation
Galectin-3	ng/mL	Fibrosis and inflammation marker

Table 5.5: *Description of Protein Markers Included in the Dataset*

Biochemical Parameter	Description
RBC (M/ μ l)	Red Blood Cell count in million per microliter
HGB (g/dL)	Hemoglobin concentration in grams per deciliter
HCT (%)	Hematocrit percentage
MCV (fl)	Mean Corpuscular Volume in femtoliters
MCH (pg)	Mean Corpuscular Hemoglobin in picograms
MCHC (g/dL)	Mean Corpuscular Hemoglobin Concentration in grams per deciliter
RDW (%)	Red Cell Distribution Width percentage
PWD (g/dL)	Plasma Width in grams per deciliter
WBC (K/ μ l)	White Blood Cell count in thousand per microliter
Neu (%)	Neutrophil percentage
Lym (%)	Lymphocyte percentage
Mon (%)	Monocyte percentage
Eos (%)	Eosinophil percentage
Bas (%)	Basophil percentage
Neu (K/ μ l)	Neutrophil count in thousand per microliter
Lym (K/ μ l)	Lymphocyte count in thousand per microliter
Mon (K/ μ l)	Monocyte count in thousand per microliter
Eos (K/ μ l)	Eosinophil count in thousand per microliter
Bas (K/ μ l)	Basophil count in thousand per microliter
Metamyelocytes + Myelocytes + Promyelocytes (%)	Percentage of immature granulocytes
Metamyelocytes + Myelocytes + Promyelocytes (K/ μ l)	Count of immature granulocytes in thousand per microliter
PLT (K/ μ l)	Platelet count in thousand per microliter
MPV (fl)	Mean Platelet Volume in femtoliters
PDW (fl)	Platelet Distribution Width in femtoliters
PCT (%)	Plateletcrit percentage
PLT (%)	Platelet percentage
UN (mg/dL)	Urea Nitrogen in milligrams per deciliter
Cr (mg/dL)	Creatinine in milligrams per deciliter
SGOT (U/L)	Aspartate Aminotransferase in units per liter
SGPT (U/L)	Alanine Aminotransferase in units per liter
c-GT (U/L)	Gamma-Glutamyl Transferase in units per liter
ALP (U/L)	Alkaline Phosphatase in units per liter
CHO (mg/dL)	Cholesterol in milligrams per deciliter
TGL (mg/dL)	Triglycerides in milligrams per deciliter
HDL (mg/dL)	High-Density Lipoprotein in milligrams per deciliter
LDL (mg/dL)	Low-Density Lipoprotein in milligrams per deciliter
GLU (mg/dL)	Glucose in milligrams per deciliter

Table 5.6: Description of Biochemical Parameters Included in the Dataset

Chapter 6

Implementation

This chapter outlines the technical implementation and theoretical framework of an end-to-end trainable system for risk stratification of carotid atheromatous plaque. It examines various innovative strategies for the integration of imaging and tabular non-image data through multimodal fusion techniques. The challenge posed by dataset imbalance is systematically mitigated, and a detailed analysis of the architecture underlying the models developed is provided. Additionally, a proof of concept is presented in the form of a web-based clinical decision support system (CDSS) prototype designed to assess plaque vulnerability.

The framework was developed using PyTorch for DL model construction and training, alongside Scikit-learn for auxiliary ML tasks. The implementation initially relied on Google Colab for GPU-accelerated prototyping but transitioned to local development due to computational constraints (credit limits and frequent re-training requirements). On macOS hardware with an AMD GPU, PyTorch's Metal Performance Shaders (MPS) backend replaced CUDA to enable hardware acceleration. An epoptive view of dataset preprocessing workflow is provided in Figure 6.1.

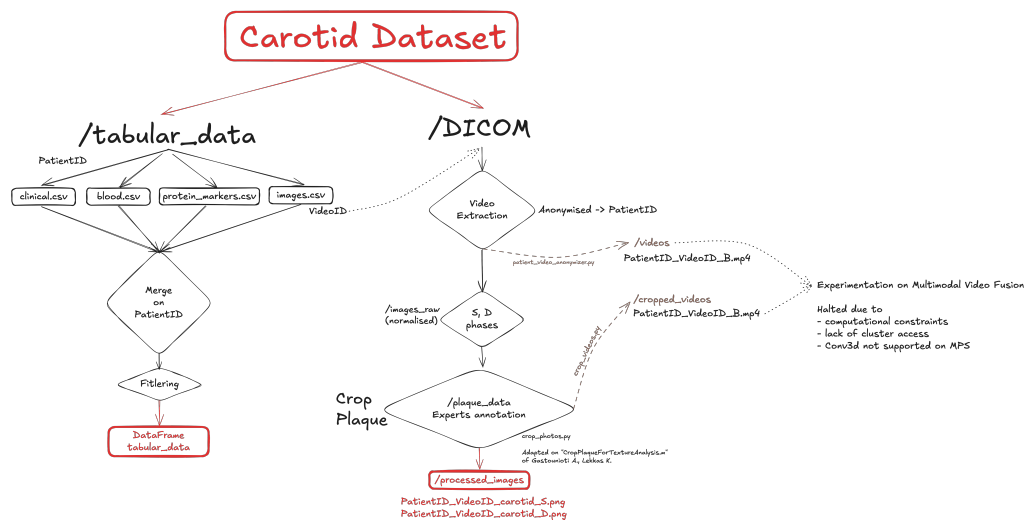


Figure 6.1: Data preprocessing workflow

6.1 Tabular data preprocessing

The pipeline begins by loading and preprocessing tabular data from structured .csv files, which are parsed into a Pandas DataFrame for subsequent manipulation. In this study, the target variable is defined as 'Risk' though alternative outcomes, such as 'Stenosis' or 'Diagnosis', could be investigated in future work.

To ensure robust model generalization, we implemented a feature selection strategy grounded in clinical relevance. Specifically, only variables with established pathophysiological significance in carotid plaque progression were retained, thereby minimizing the risk of overfitting to spurious correlations. Nonessential feature, derived from clinical, biochemical, and protein marker datasets were rigorously filtered based on their limited association with vulnerable plaque development according to latest available literature. The selected features were then integrated into a unified DataFrame, forming the structured input for subsequent modeling. For traceability and multimodal integration, patient and video identifiers were extracted to facilitate later alignment with imaging data. Categorical variable 'Smoker' undergoes one-hot encoding generating Current Smoker and Old Smoker features, according to Table 6.1. A summary of tabular data features used as input in the models are presented in Table 6.2 and a detailed description in Table 6.3.

Smoker	Current Smoker	Old Smoker
1	0	0
2	0	1
3	1	0

Table 6.1: *One-Hot Encoding of the 'Smoker' Feature*

Category	Parameters
Clinical	Gender, Age, Old Smoker, Current Smoker, Diabetes, Hypertension, Anti-diabetics, Antihypertensives, Dyslipidemia, CAD
Biochemical	PLT (K/ μ L), UN (mg/dL), Cr (mg/dL), SGOT/AST (U/L), SGPT/ALT (U/L), γ -GT (U/L), ALP (U/L), CHO (mg/dL), TGL (mg/dL), HDL (mg/dL), LDL (mg/dL), GLU (mg/dL)
Protein Biomarkers	CRP (ng/mL), Fibrinogen (pg/mL), MMP-1, MMP-2, MMP-7, MMP-9, IL-1 β , IL-6, TNF- α (pg/mL), TIMP-1, TIMP-2 (pg/mL), C-peptide, Insulin, RBP4 (pg/mL), Galectin-3 (ng/mL)

Table 6.2: *Summary of Tabular Data Features Used as Input in the Model*
All matrix metalloproteinases (MMP), tissue inhibitors of metalloproteinases (TIMP), and interleukins (IL) measured in pg/mL unless noted.

6.2 Imaging data preprocessing

The image preprocessing pipeline in our implementation is designed to ensure standardized, high-quality inputs for the deep learning models. To maintain compatibility with common deep learning libraries, images are first converted from OpenCV's default BGR format to the RGB color space. Each image is then resized to a target resolution of 224×224 pixels, aligning with the input dimensions expected by popular CNNs such as ResNet. The preprocessing workflow includes the following steps:

1. An optimal scaling factor is computed to resize the image while preserving the original aspect ratio, ensuring no distortion occurs.
2. Resizing is performed using Lanczos interpolation, chosen for its high-quality anti-aliasing properties, critical for preserving subtle anatomical details.

The Lanczos interpolation kernel is defined as a windowed sinc function, which helps in high-quality image resizing while minimizing aliasing artifacts. The one-dimensional Lanczos kernel of order a is given by:

$$L(x) = \begin{cases} \text{sinc}(x) \cdot \text{sinc}\left(\frac{x}{a}\right), & \text{if } |x| < a, \\ 0, & \text{otherwise,} \end{cases} \quad (6.1)$$


where $\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$ (with $\text{sinc}(0) = 1$), and a is a positive integer (typically $a = 2$ or 3).

For two-dimensional image resizing, the interpolation is applied separably along each axis. The resampled pixel value $I'(x, y)$ at a new coordinate (x, y) is computed as a weighted sum of neighboring pixels in the original image I :

$$I'(x, y) = \sum_{i=\lfloor x \rfloor - a + 1}^{\lfloor x \rfloor + a} \sum_{j=\lfloor y \rfloor - a + 1}^{\lfloor y \rfloor + a} I(i, j) \cdot L(x - i) \cdot L(y - j), \quad (6.2)$$

where $\lfloor \cdot \rfloor$ denotes the floor function, and the kernel $L(\cdot)$ ensures smooth transitions while preserving high-frequency details.

3. After resizing, black (zero-value) padding is applied to center the image within the 224×224 frame. The exact padding applied on each side (left, right, top, bottom) is recorded for potential use in future postprocessing or image reconstruction tasks.
4. Finally, pixel values are normalized to the $[0, 1]$ range by dividing by 255. However, since images were pre-normalized, this step is omitted in practice.

Figure 6.2: *Original and processed image data samples*

6.3 Training and model evaluation

The dataset comprises 65 high-risk and 19 low-risk carotid plaque cases, exhibiting a pronounced class imbalance. Notably, three patients have imaging records from both RICA and LICA, with each side classified under different risk categories. Additionally, six patients contribute two low-risk imaging records each, while two patients provide multiple high-risk imaging records, as depicted in Figure 6.3. This distribution introduces two critical challenges:

1. the scarcity of low-risk samples risks biasing model training toward majority-class patterns, particularly under conventional data partitioning which may allocate insufficient low-risk samples to validation and test sets
2. patient-level data interdependence, where multiple images from the same individual span training and evaluation splits threatens to artificially inflate performance metrics due to dataset leakage.

The training process employs a stratified group k-fold cross-validation ($k=3$) to handle limited data while ensuring balanced class distribution and preventing patient data leakage across splits. To mitigate the class imbalance a Weighted Random Sampling is applied along with a Weighted Loss Function. For each fold, the model is trained with




Figure 6.3: *Patients records distribution in Low and High Risk classes*

data augmentation specifically designed for ultrasound images, including minimal spatial transformations and intensity adjustments. By using image specific augmentations (preserving modality alignment) coupled with loss reweighting, the approach maintains dataset integrity while addressing imbalance. This avoids the need for post-hoc validation of synthetic sample plausibility, a non trivial task requiring domain expertise to verify clinical, biochemical, protein biomarkers and imaging correlations for each synthetic instance.

We employ a series of transformations designed to improve model generalization while preserving clinically relevant features in carotid ultrasound images. Spatial augmentations include random horizontal flips ($p = 0.5$) and subtle affine transformations with minimal rotation ($\pm 2^\circ$), tiny translations ($\pm 2^\circ$), and minor scaling (0.98 - 1.02) to simulate probe variability. Intensity-based augmentations incorporate micro-rotations ($p = 0.3$), brightness/contrast adjustments ($\pm 10^\circ$), and Gaussian noise ($\sigma = 0.03$) to mimic ultrasound acquisition artifacts. A summary of set values of parameters is provide at Table 6.4. These transformations are applied sequentially, with careful value clipping to maintain valid pixel ranges. The validation and test sets remain unaugmented to ensure unbiased evaluation.

Models performance is evaluated using several complementary metrics beyond accuracy, including sensitivity, specificity, F1-score, MCC score, and area under the receiver operating characteristic curve (AUC-ROC), as defined in 3.4.

6.4 Models Architecture

The multifactorial pathophysiology of carotid atherosclerosis, as covered in depth in Chapter 2, spanning localized hemodynamic stresses, systemic inflammatory cascades, and metabolic dysregulation demands architectures capable of modeling hierarchical interactions between imaging derived properties and clinical/biological context. Traditional late fusion strategies, which process modalities independently before combining predictions, fail to capture the intrinsic coupling between plaque morphology and circulating biomarkers. Conversely, attention-based joint fusion enables dynamic, context-aware integration where learned attention weights quantify the relative contribution of imaging versus clinical features at the feature level, mirroring the clinical reality where plaque vulnerability emerges from neither modality in isolation but their synergistic interplay. This approach aligns with the governing equation of plaque vulnerability, where

$$\text{Plaque Vulnerability} = f \left(\underbrace{\text{Local Hemodynamics}}_{\text{Imaging}}, \underbrace{\text{Systemic Inflammation}}_{\text{Biomarkers}}, \underbrace{\text{Metabolic State}}_{\text{Clinical}} \right) \quad (6.3)$$

Cross-modal attention mechanisms further refine this integration by explicitly modeling interactions between ultrasound derived texture features, such as necrotic core echogenicity, and proteomic signatures. Early fusion was explored as a baseline, concatenating raw image pixels with normalized tabular data, though its effectiveness is inherently limited by the disparate dimensional and semantic nature of these inputs. Hybrid architectures, including Attention-Gated Video Hybrid Fusion, were designed to address spatiotemporal dependencies in ultrasound cine loops while preserving clinical context, though computational constraints prioritized 2D implementations. Collectively, these architectures operationalize the biological premise that plaque instability arises from concurrent local mechanical failure and systemic molecular dysregulation, necessitating models that learn *how*, *not just whether*, modalities interact.

6.4.1 Joint Attention Based Fusion

We implement a joint attention based fusion strategy that dynamically combines deep representations from imaging and tabular data using an attention mechanism. By concatenating modality-specific features and learning adaptive attention weights, this approach selectively emphasizes the most discriminative features from each modality before classification.

CNN Backbones used for image feature extraction

For our experiments, we selected three CNN backbones pre-trained on ImageNet, employing a transfer learning approach. Specifically, we used the following architectures:

1. ResNet18, architecture presented in Figure 6.4
2. VGG16, architecture presented in Figure 6.5

3. EfficientNet-B0, architecture presented in Figure 6.6

This decision was driven by the constraints imposed by the relatively small size of our dataset, which limits the feasibility of using larger, more complex architectures. Given the limited data, utilizing such models would likely lead to overfitting unless substantial regularization techniques were applied. To mitigate this risk, we opted to freeze all layers of the pre-trained models except for the final parameters. This strategy ensures that the majority of the model's learned features are preserved, while the last layers, which are more directly responsible for the task specific predictions, are fine-tuned for our specific dataset. In doing so, we achieve a balance between leveraging the pre-learned representations from ImageNet and adapting the model to the particularities of our experimental context. Consequently, the selected models strike a balance between performance and the risk of overfitting, making them more suitable for our experimental setup.




Figure 6.4: *ResNet18 Architecture*
[89]

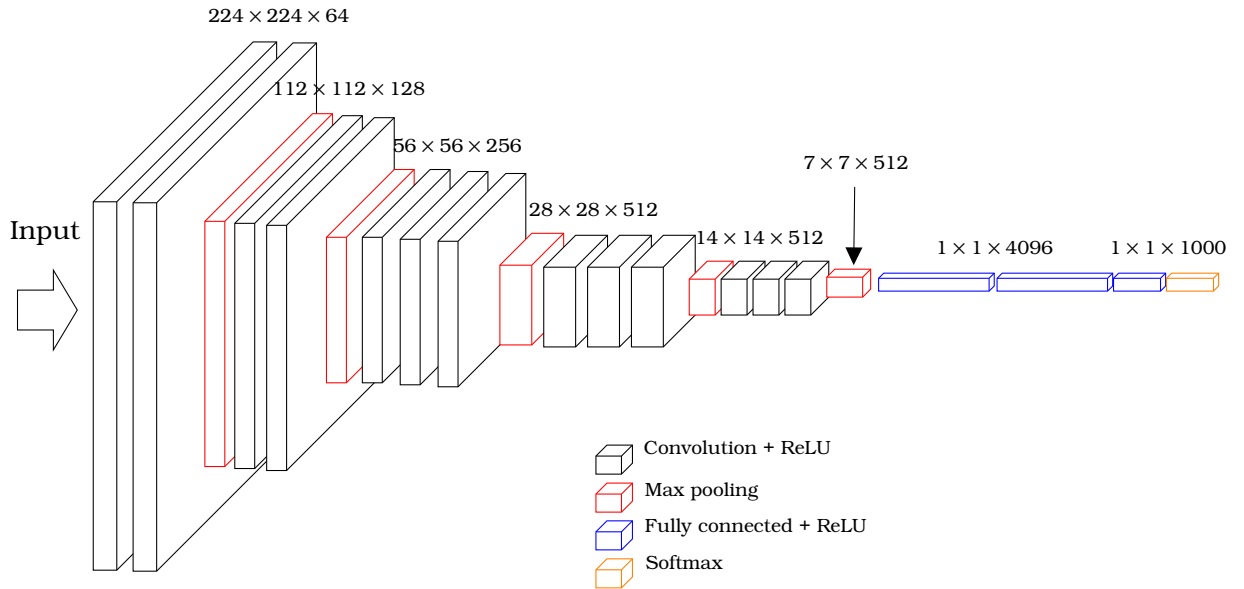


Figure 6.5: *VGG16 Architecture*

The Multimodal Joint Attention-based Fusion model operates through three main phases, each designed to progressively integrate and refine information from each input modality. The overall architecture of the model is depicted in Figure 6.7.

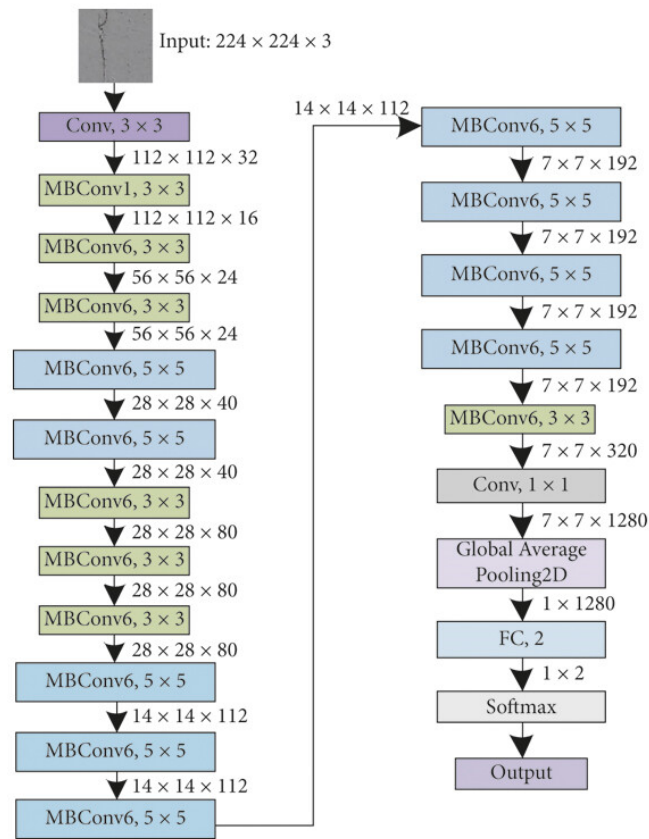


Figure 6.6: *EfficientNet-B0 Architecture*
[90]

Phase 1: Modality-Specific Feature Extraction

Image Pathway

- **Backbone:** Pretrained CNN (ResNet18, VGG16, EfficientNet)
- **Input:** Stacked systolic + diastolic ultrasound images $(B, 2, 3, H, W)$
- **Processing:**
 - Extracts 128-dimensional features per image
 - Temporal pooling: Averages features across cardiac phases
- **Output:** $(B, 128)$ image embedding

Tabular Pathway

- **Backbone:** 2-Layer MLP (ReLU, BatchNorm, Dropout)
- **Input:** Clinical/demographic features $(B, \text{tabular_dim})$
- **Processing:**
 - Linear projections: $\text{tabular_dim} \rightarrow 64 \rightarrow 32$
- **Output:** $(B, 32)$ tabular embedding

Phase 2: Attention-Based Fusion

Concatenation

- Combines image and tabular embeddings into a unified representation: $(B, 160)$

Attention Gate

- **Architecture:** Linear($160 \rightarrow 2$) followed by Softmax
- **Dynamically predicts weights for:**
 - a_{img} (image importance)
 - a_{tab} (tabular importance)

Weighted Fusion

- Re-scales features based on their importance:
 - $a_{\text{img}} \cdot \mathbf{h}_{\text{img}}$
 - $a_{\text{tab}} \cdot \mathbf{h}_{\text{tab}}$
- **Final representation:** $(B, 160)$

Phase 3: Classification

Head Architecture

- Linear($160 \rightarrow 64$) with BatchNorm, ReLU, and Dropout
- Linear($64 \rightarrow \text{num_classes}$)

Output

- Logits for classification into target classes
- Attention weights for interpretability

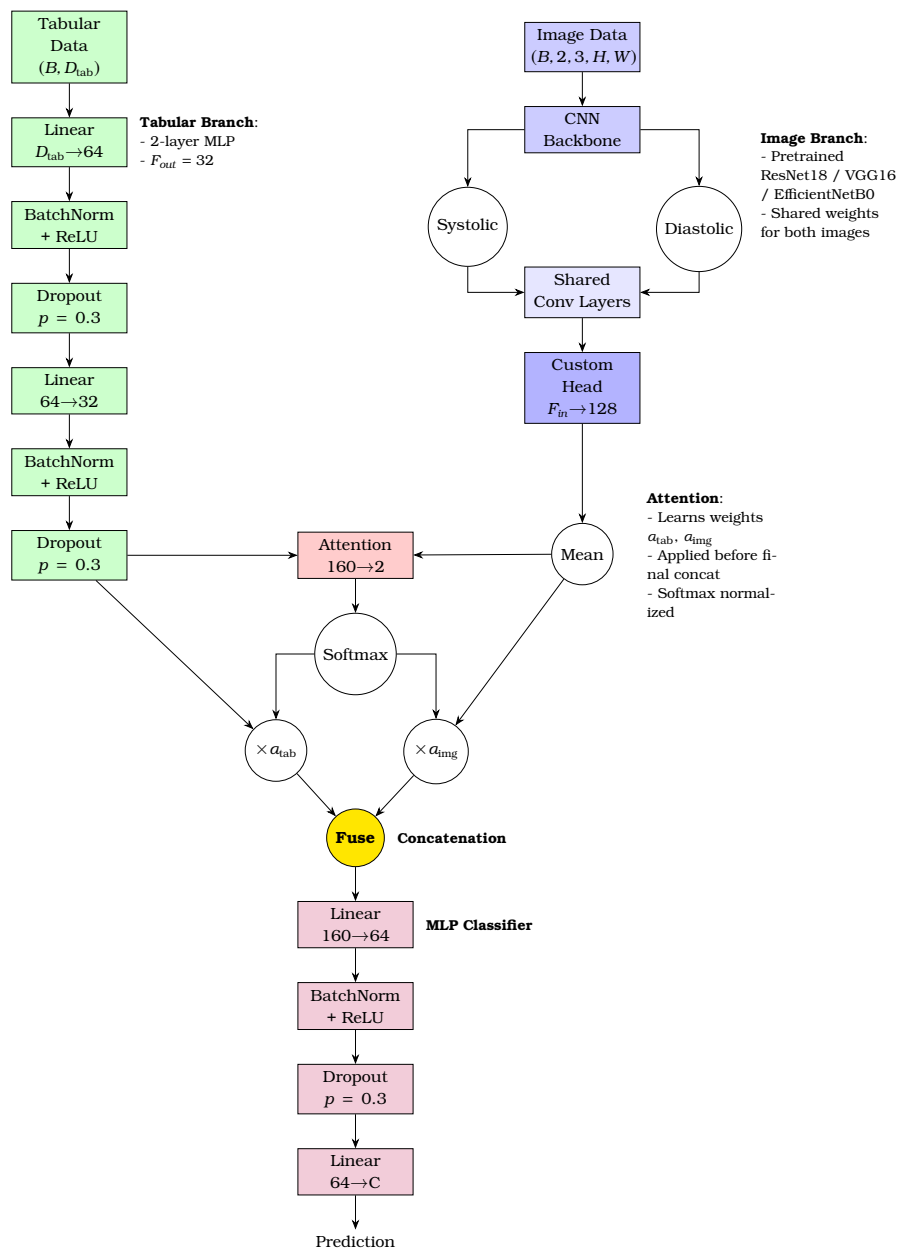


Figure 6.7: Joint Attention Based Fusion model architecture
 Color coding: tabular processing, image processing, fusion, attention, classifier.

6.4.2 Early fusion

The early fusion architecture (Figure 6.8) serves as a foundational baseline for integrating systolic and diastolic ultrasound images with clinical tabular data through concatenation at the feature level. This approach processes ultrasound frames via dual EfficientNet-B0 backbones, one for each cardiac phase, to extract high-dimensional image embeddings (B,1280), while a parallel MLP encodes tabular features into a compact 64-dimensional vector. Concatenation of these modality-specific representations creates a unified feature space (B,2624), subsequently refined through a cascade of dense layers with batch normalization and dropout.

While conceptually straightforward, this architecture inherently assumes linear compatibility between imaging and clinical subspaces, a tenuous premise given the biological reality where plaque vulnerability emerges from nonlinear interactions between local tissue mechanics, such as plaque strain patterns, and systemic factors as CRP and IL-6-driven inflammation.

The model's rigidity becomes apparent in its static treatment of modality contributions. Concatenation assigns equal a priori importance to all features, disregarding context dependent relevance. It would not be able for example to prioritize lipid-core texture descriptors when CRP levels exceed clinical thresholds. Nevertheless, early fusion provides a critical comparative framework, demonstrating how naive fusion strategies struggle to reconcile the disparate scales and semantics of ultrasound-derived spatiotemporal features (dimensions) with low-dimensional clinical variables.

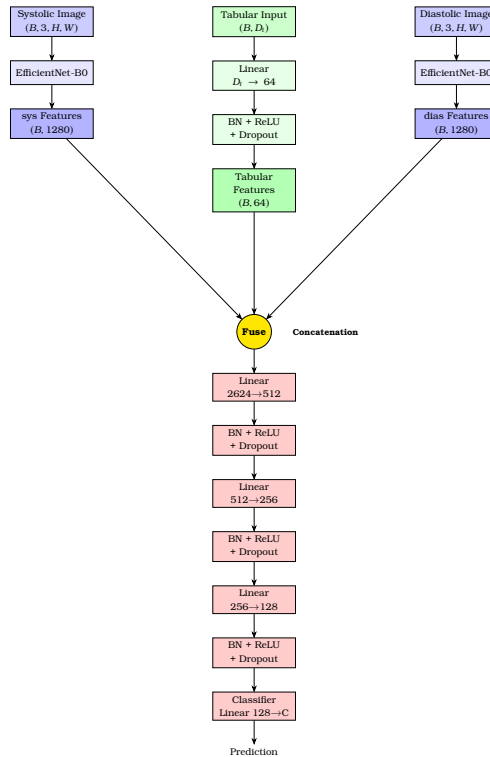


Figure 6.8: *Early Fusion model architecture*
Color coding: image processing, tabular processing, fusion, classifier.

6.4.3 Attention-Gated Video Hybrid Fusion

Preliminary experiments on multimodal video fusion were conducted to integrate kinematic, hemodynamic, and textural data streams for enhanced plaque vulnerability assessment. The proposed architecture leveraged 3D spatiotemporal convolutions alongside attention-based multimodal fusion, demonstrating feasibility in feature extraction from ultrasound sequences. However, scaling these experiments proved prohibitively resource intensive due to

1. $O(T \times H \times W)$ memory complexity of volumetric convolutions
2. hardware limitations in Apple's Metal Performance Shaders (MPS) framework for 3D convolution operations
3. insufficient GPU memory for batch processing of high-resolution cine loops

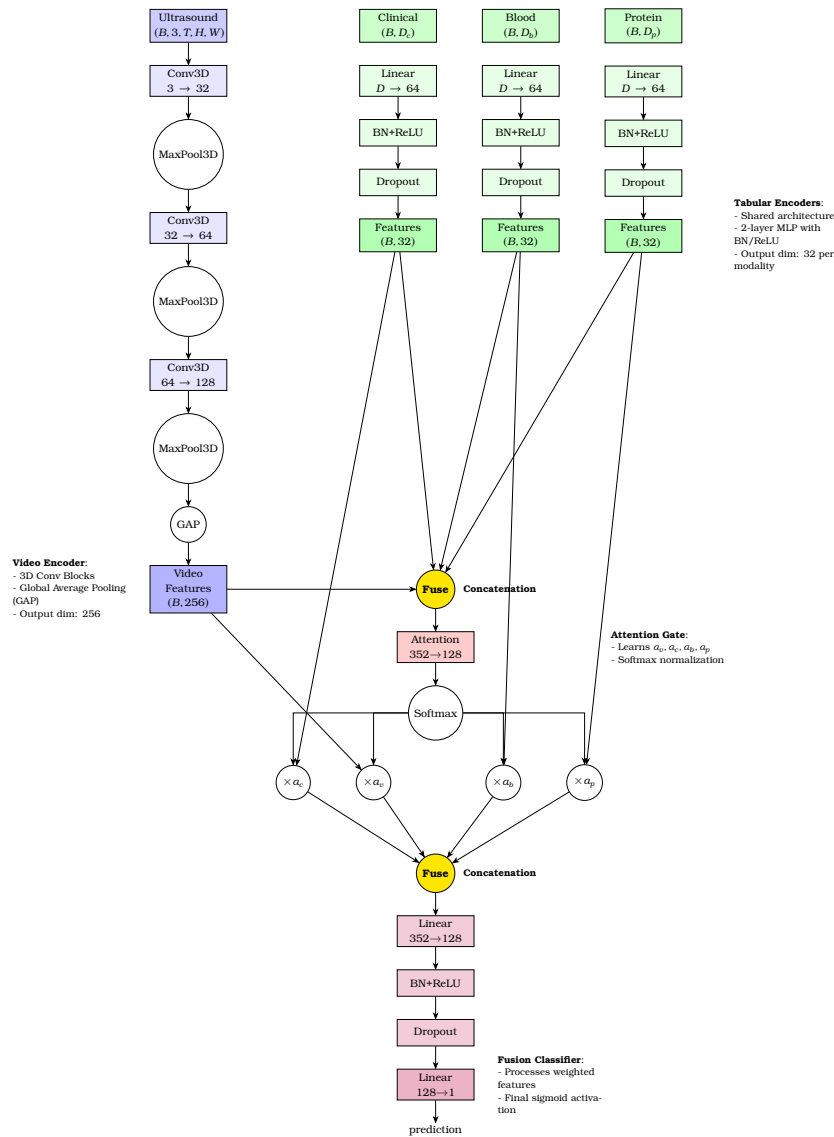


Figure 6.9: Architecture of Attention-Gated Video Hybrid Fusion
Color coding: video processing, tabular processing, fusion, attention, classifier.

6.5 A prototype CDSS for plaque vulnerability assessment

In the context of this work, an initial version of a Clinical Decision Support System (CDSS) for the assessment of carotid plaque vulnerability was implemented as a web-based prototype, illustrated in Figure 6.10, that leverages multimodal data fusion of ultrasound imaging and patient-specific clinical information. The system integrates systolic and diastolic B-mode carotid ultrasound images with demographic, laboratory, and biomarker data to provide a real-time risk score predicting plaque instability using a TensorFlow.js DL model deployed entirely in the browser. The initial experimental web application architecture is shown in Figure 6.11, with all processing taking place into browser's JavaScript engine to ensure medical data protection.

The screenshot displays the 'Carotid Plaque Vulnerability Assessment' Clinical Decision Support System (CDSS) interface. The header is a blue bar with the title and subtitle. The main content is divided into two panels. The left panel, 'Ultrasound Images', shows the 'Systolic Phase' with a file upload button and a selected file 'US_B_Mode_Diastolic'. Below it is a B-mode ultrasound image of a carotid artery with a 'Crop Image' button. The 'Diastolic Phase' section below it has a 'Choose file' button and a 'Crop Image' button. The right panel, 'Clinical Data', features a progress bar at 44% and tabs for 'Demographics', 'Medical History', 'Lab Values', and 'Biomarkers'. The 'Demographics' tab is active, showing input fields for 'Age' (71), 'Gender' (Male), 'Current Smoker' (Yes), and 'Former Smoker' (Select). A green 'Analyze Plaque Vulnerability' button is at the bottom of the clinical data section. A yellow disclaimer banner at the bottom states: 'This tool is for research use only. It is not approved for clinical or diagnostic use. Any risk assessments or predictions generated by this system are experimental and must not be used as the sole basis for medical decisions. Clinicians must rely on their professional expertise and a complete set of patient data when making any diagnostic or therapeutic decisions. v0.1 - Alpha'.

Figure 6.10: CDSS prototype for the assessment of carotid plaque vulnerability

While this tool currently serves as a proof-of-concept, its design lays the foundation for future clinical application pending thorough evaluation of model performance and generalizability across diverse patient populations. If such validation confirms robust predictive capabilities and external reproducibility, the system could even be considered for clinical assistance especially if future iterations enhance its interpretability and explainability, allowing clinicians to understand not just the risk score but also the contributing features behind each prediction.

Such transparency is critical for building trust in AI based CDSS, particularly in high-stakes medical decisions. A clinically interpretable system could provide visual and textual rationales aligned with recognized plaque features or biomarker profiles, empowering clinicians to correlate model outputs with their own observations. This would be particularly valuable for vascular surgeons, neurologists, and interventional radiologists

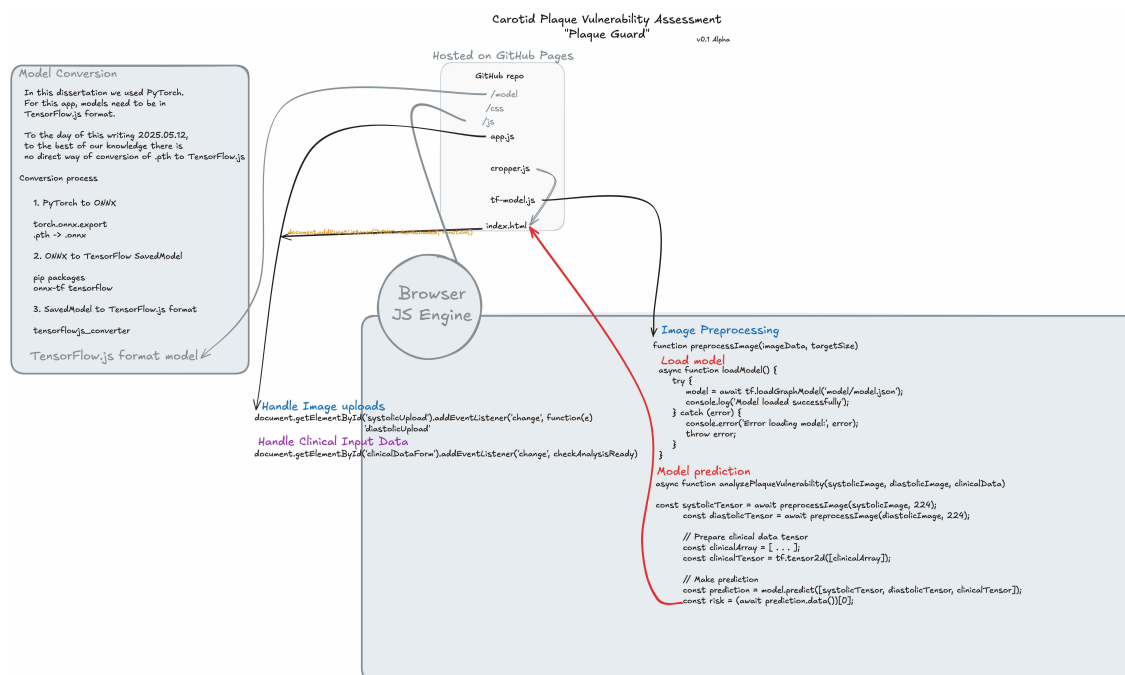


Figure 6.11: CDSS architecture

involved in the management of carotid artery disease, where treatment decisions such as surgical intervention versus conservative therapy depend on comprehensive assessments of plaque stability. By offering an integrated view of multimodal patient data and linking it to established guidelines, the tool could enhance decision-making precision, reduce variability, and ultimately support better patient outcomes in real-world settings.

Furthermore, the tool is intended to evolve in alignment with established clinical frameworks. It will incorporate recommendations from the most recent European Society for Vascular Surgery (ESVS) Clinical Practice Guidelines on the Management of Atherosclerotic Carotid and Vertebral Artery Disease (Supplementary Table A.1), as well as integrate the American Heart Association (AHA) atherosclerotic plaque classification system (Supplementary Table A.3). This structured integration will ensure that the CDSS not only reflects the latest evidence-based practices but also provides actionable guidance tailored to individual plaque characteristics, potentially supporting more precise therapeutic decision-making in the future.

Parameter	Description
Clinical Parameters	
Gender	Male or female sex.
Age	Age in years.
Old Smoker	History of past smoking (binary: yes/no).
Current Smoker	Active smoking status (binary: yes/no).
Diabetes	Presence of diabetes mellitus (binary: yes/no).
Hypertension	Presence of high blood pressure (binary: yes/no).
Anti-diabetics	Use of diabetes medications prior admission (binary: yes/no).
Antihypertensives	Use of antihypertensives drugs prior admission (binary: yes/no).
Dyslipidemia	Abnormal lipid metabolism (binary: yes/no).
CAD	Coronary Artery Disease (binary: yes/no).
Biochemical Parameters	
PLT (K/ μ L)	Platelet count (thousands per microliter).
UN (mg/dL)	Urea nitrogen level in blood.
Cr (mg/dL)	Serum creatinine, a kidney function marker.
SGOT/AST (U/L)	Aspartate aminotransferase, liver/heart enzyme.
SGPT/ALT (U/L)	Alanine aminotransferase, liver enzyme.
γ -GT (U/L)	Gamma-glutamyl transferase, liver/biliary marker.
ALP (U/L)	Alkaline phosphatase, liver/bone enzyme.
CHO (mg/dL)	Total cholesterol level.
TGL (mg/dL)	Triglyceride level.
HDL (mg/dL)	High-density lipoprotein.
LDL (mg/dL)	Low-density lipoprotein.
GLU (mg/dL)	Blood glucose level.
Protein Biomarkers	
CRP (ng/mL)	C-reactive protein, inflammation marker.
Fibrinogen (pg/mL)	Blood clotting factor.
MMP-1, -2, -7, -9 (pg/mL)	Matrix metalloproteinases (tissue remodeling enzymes).
IL-1 β , IL-6 (pg/mL)	Pro-inflammatory interleukins.
TNF- α (pg/mL)	Tumor necrosis factor-alpha, inflammatory cytokine.
TIMP-1, -2 (pg/mL)	Tissue inhibitors of metalloproteinases.
C-peptide (pg/mL)	Marker of insulin production.
Insulin (pg/mL)	Hormone regulating blood glucose.
RBP4 (pg/mL)	Retinol-binding protein 4, linked to insulin resistance.
Galectin-3 (ng/mL)	Protein involved in fibrosis and inflammation.

Table 6.3: *Detailed Description of Clinical, Biochemical, and Protein Biomarkers Used as Input in the Implemented Models*

Abbreviations: CAD = Coronary artery disease; PLT = Platelets; UN = Urea nitrogen; Cr = Creatinine; SGOT/AST = Serum glutamic-oxaloacetic transaminase/Aspartate transaminase; SGPT/ALT = Serum glutamic-pyruvic transaminase/Alanine transaminase; γ -GT = Gamma-glutamyl transferase; ALP = Alkaline phosphatase; CHO = Cholesterol; TGL = Triglycerides; HDL/LDL = High-/Low-density lipoprotein; GLU = Glucose; CRP = C-reactive protein; MMP = Matrix metalloproteinase; IL = Interleukin; TNF- α = Tumor necrosis factor-alpha; TIMP = Tissue inhibitor of metalloproteinases; RBP4 = Retinol-binding protein 4.

Parameter	Value
Random Horizontal Flip Probability	0.5
Rotation Range	$\pm 2^\circ$
Translation Range	± 0.2
Scale Range	0.98-1.02
Brightness Adjustment	± 0.1
Contrast Adjustment	± 0.1
Noise Level	0.3
Micro-rotation Probability	0.3

Table 6.4: *Image Augmentation Parameters*

Results

This chapter presents the comprehensive experimental results of our deep multi-modal fusion approaches. We begin by detailing the performance metrics of our Joint Attention Based Fusion models employing three distinct CNN backbones (ResNet18, VGG16, and EfficientNet), with particular focus on their classification performance and the learned attention weights between imaging and tabular data modalities. The chapter further explores the effect of imposing various thresholds on image attention weights to achieve optimal modality balance. Through an ablation study, we investigate the contribution of each modality to the overall performance of our multimodal architecture. Additionally, we present findings from our Early Fusion model implementation to provide comparative insights into different fusion strategies. All experiments were conducted using standardized training parameters and cross-validation protocols as outlined in the methodology, with results reported across multiple performance metrics to enable robust evaluation of each approach’s clinical utility and statistical significance.

For the training scheme of our implemented models we employed adaptive learning rate optimization (Adam) with an initial learning rate of 1e-3 and early stopping criteria monitoring validation loss with a patience of 10 epochs. To address class imbalance during training, class weights were applied proportionally to the inverse frequency of each class, thereby increasing the contribution of minority class samples to the loss function. A summary of set training parameters is provided at Table 7.1

Parameter	Value
Number of Folds	3
Batch Size	4-8 (adaptive)
Optimizer	Adam
Learning Rate	0.001
Weight Decay (L2)	1e-4
Epochs	50
Early Stopping Patience	10
Loss Function	Weighted Cross-Entropy
Class Weighting	Inverse Frequency
Scheduler	ReduceLROnPlateau
Reduction Factor	0.5
Patience	3

Table 7.1: Summary of set training parameters

7.1 Joint Attention Based Fusion Models

The Joint Attention Based Fusion models demonstrate promising performance across three different CNN backbone architectures, with the EfficientNet backbone achieving the highest area under the curve (AUC) of $86.07 \pm 6.96\%$ and specificity of $64.88 \pm 7.48\%$. The plots of metrics across the fold for the best performed model are illustrated in Figure 7.2, while the confusion matrices are shown in Figure 7.1. The ResNet18 implementation delivered the highest overall accuracy ($79.70 \pm 2.19\%$) and sensitivity ($89.72 \pm 8.91\%$), indicating strong performance in detecting high-risk plaques. However, its specificity was relatively low ($52.38 \pm 20.69\%$), indicating limited effectiveness in detecting low-risk plaques.

A particularly illuminating finding across all model implementations was the naturally learned attention weight distribution, which consistently assigned substantially higher importance to imaging data (69.4%-76.1%) compared to tabular clinical data (23.9%-30.6%), suggesting ultrasound images contain more discriminative features for plaque risk assessment. The elevated standard deviations in performance metrics stem from the small test set size, which increases susceptibility to variability during fold-wise cross-validation.

When various thresholds were applied to limit image attention weights and increase tabular data contribution (as shown in Tables 7.3, 7.4, and 7.5), performance metrics generally decreased, confirming that the naturally learned attention distribution optimizes classification performance. The VGG16 backbone consistently underperformed compared to other architectures across all metrics, indicating its feature extraction capabilities may be less suited for this particular task and that further optimizations are needed.

Metric	ResNet18	VGG16	EfficientNet
Accuracy	79.70 ± 2.19	68.81 ± 9.38	77.19 ± 6.65
Balanced Accuracy	71.05 ± 6.14	62.36 ± 1.85	73.28 ± 2.18
AUC	78.15 ± 5.09	67.02 ± 9.94	86.07 ± 6.96
Specificity	52.38 ± 20.69	47.02 ± 20.12	64.88 ± 7.48
Sensitivity	89.72 ± 8.91	77.69 ± 18.05	81.67 ± 10.64
Precision	83.94 ± 1.74	75.87 ± 4.20	82.11 ± 4.33
F1 Score	78.75 ± 3.76	69.35 ± 5.39	78.42 ± 4.92
MCC	0.454 ± 0.040	0.253 ± 0.054	0.445 ± 0.112
a_{tabular}	0.239	0.302	0.306
a_{image}	0.761	0.698	0.694

Table 7.2: *Performance Metrics of Joint Fusion Attention-Based Models*

Indicated metrics: $\overline{metric} \pm SD$ [0, 100%] on stratified 3-fold cross-validation

MCC: $\overline{mcc} \pm SD$ [-1, 1]

7.2 Ablation study

To better understand the contribution of each modality in the Joint Attention-Based Fusion Model we conducted an ablation study, on the best performed architecture (Joint

Metric	No threshold	$a_{image} \leq 0.55$	$a_{image} \leq 0.65$	$a_{image} \leq 0.7$
Accuracy	79.70 ± 2.19	76.13 ± 4.81	76.17 ± 1.85	73.53 ± 9.61
Balanced Accuracy	71.05 ± 6.14	63.68 ± 9.78	66.07 ± 11.37	64.37 ± 7.64
AUC	78.15 ± 5.09	71.80 ± 7.82	77.33 ± 3.32	73.15 ± 8.37
Specificity	52.38 ± 20.69	37.50 ± 30.62	45.83 ± 32.81	43.45 ± 25.52
Sensitivity	89.72 ± 8.91	89.86 ± 14.35	86.30 ± 10.69	85.29 ± 17.39
Precision	83.94 ± 1.74	75.96 ± 13.03	73.84 ± 11.85	79.34 ± 9.33
F1 Score	78.75 ± 3.76	73.01 ± 5.66	73.83 ± 5.94	72.36 ± 8.38
MCC	0.454 ± 0.040	0.292 ± 0.225	0.281 ± 0.204	0.330 ± 0.199
$a_{tabular}$	0.239	0.460	0.382	0.340
a_{image}	0.761	0.540	0.618	0.660

Table 7.3: *Performance Metrics of the Joint Fusion Attention Model with ResNet-18 Backbone at Varying Image Attention Weight Thresholds*

Indicated metrics: $\overline{metric} \pm SD$ [0, 100%] on stratified 3-fold cross-validation.

MCC: $\overline{mcc} \pm SD$ [-1, 1]

Metric	No Threshold	$a_{image} \leq 0.55$	$a_{image} \leq 0.65$
Accuracy	68.81 ± 9.38	67.49 ± 12.62	68.68 ± 12.89
Balanced Accuracy	62.36 ± 1.85	60.23 ± 3.88	64.13 ± 9.72
AUC	67.02 ± 9.94	71.61 ± 1.84	78.41 ± 8.12
Specificity	47.02 ± 20.12	42.86 ± 22.78	52.38 ± 20.69
Sensitivity	77.69 ± 18.05	77.61 ± 21.93	75.88 ± 17.91
Precision	75.87 ± 4.20	77.07 ± 9.30	76.32 ± 10.62
F1 Score	69.35 ± 5.39	67.45 ± 8.56	69.89 ± 10.94
MCC	0.253 ± 0.054	0.255 ± 0.162	0.281 ± 0.229
$a_{tabular}$	0.302	0.450	0.425
a_{image}	0.698	0.550	0.575

Table 7.4: *Performance Metrics of the Joint Fusion Attention Model with VGG-16 Backbone at Varying Image Attention Weight Thresholds*

Indicated metrics: $\overline{metric} \pm SD$ [0, 100%] on stratified 3-fold cross-validation

MCC: $\overline{mcc} \pm SD$ [-1, 1]

Attention-Based Fusion Model with EfficientNetB0 backbone). Specifically, we isolated the image and tabular branches by evaluating them independently while keeping the rest of the training and evaluation pipeline unchanged, as illustrated schematically in Figure 7.3.

In the first experiment, we retained only the image branch of the model, removing the tabular input while preserving the architecture, preprocessing, and training settings. This allowed us to assess the standalone performance of the imaging modality.

In the second experiment, we did the opposite: we kept only the tabular branch and removed the image input. Again, all other components of the pipeline were maintained to ensure a fair comparison.

The multimodal model achieved an accuracy of $77.19 \pm 6.65\%$ and AUC of $86.07 \pm 6.96\%$, compared to $74.85 \pm 5.79\%$ accuracy and $84.55 \pm 7.91\%$ AUC with imaging data alone, and considerably lower performance with tabular data alone ($71.62 \pm 7.04\%$ accuracy, $64.61 \pm 6.65\%$ AUC). This performance differential is further emphasized by the Matthews Correlation Coefficient (MCC), which measures overall classification quality,

Metric	No Threshold	$a_{image} \leq 0.55$	$a_{image} \leq 0.65$
Accuracy	77.19 ± 6.65	76.13 ± 4.81	70.08 ± 5.35
Balanced Accuracy	73.28 ± 2.18	59.29 ± 12.01	58.79 ± 7.70
AUC	86.07 ± 6.96	74.09 ± 8.61	73.09 ± 11.32
Specificity	64.88 ± 7.48	29.17 ± 25.69	37.50 ± 27.00
Sensitivity	81.67 ± 10.64	89.42 ± 8.88	80.07 ± 14.65
Precision	82.11 ± 4.33	71.56 ± 10.22	69.83 ± 8.86
F1 Score	78.42 ± 4.92	73.29 ± 6.68	68.61 ± 2.80
MCC	0.445 ± 0.112	0.192 ± 0.252	0.155 ± 0.148
$a_{tabular}$	0.306	0.452	0.409
a_{image}	0.694	0.548	0.591

Table 7.5: *Performance Metrics of the Joint Fusion Attention Model with EfficientNet Backbone at Varying Image Attention Weight Thresholds*

Indicated metrics: $\overline{metric} \pm SD$ [0, 100%] on stratified 3-fold cross-validation

MCC: $\overline{mcc} \pm SD$ [-1, 1]

where the multimodal approach scored 0.445 ± 0.112 versus 0.355 ± 0.063 for imaging-only and 0.257 ± 0.076 for tabular-only approaches.

The relatively strong performance of the imaging-only model confirms the rich discriminative information content of ultrasound images, while the significant performance boost when combining modalities validates the fundamental premise of multimodal fusion, that complementary information streams can enhance classification performance beyond what is possible with any single modality. Notably, the balanced accuracy metric shows the most substantial improvement in the multimodal approach ($73.28 \pm 2.18\%$) compared to either imaging ($65.75 \pm 6.09\%$) or tabular ($63.60 \pm 4.88\%$) modalities alone, indicating better performance across both positive and negative classes.

Metric	Both Modalities	Imaging Modality	Tabular Modality
Accuracy	77.19 ± 6.65	74.85 ± 5.79	71.62 ± 7.04
Balanced Accuracy	73.28 ± 2.18	65.75 ± 6.09	63.60 ± 4.88
AUC	86.07 ± 6.96	84.55 ± 7.91	64.61 ± 6.65
Specificity	64.88 ± 7.48	46.43 ± 24.91	52.38 ± 20.69
Sensitivity	81.67 ± 10.64	85.07 ± 14.62	74.83 ± 11.72
Precision	82.11 ± 4.33	81.71 ± 2.99	75.52 ± 3.07
F1 Score	78.42 ± 4.92	73.93 ± 2.62	72.46 ± 5.63
MCC	0.445 ± 0.112	0.355 ± 0.063	0.257 ± 0.076

Table 7.6: *Ablation Study Findings*

Indicated metrics: $\overline{metric} \pm SD$ [0, 100%] on stratified 3-fold cross-validation

MCC: $\overline{mcc} \pm SD$ [-1, 1]

7.3 Early Fusion Model

The Early Fusion model, implemented using dual EfficientNet-B0 architectures for concurrent processing of imaging and non-imaging modalities, demonstrated moderate but suboptimal performance characteristics when evaluated against the attention-based multimodal fusion approach. The model achieved an overall accuracy of $76.08 \pm 4.05\%$

with a corresponding Area Under the Curve (AUC) of $74.37 \pm 14.19\%$, indicating reasonable discriminative capability but with notable performance variability across validation folds.

A detailed examination of the classification metrics reveals an inherent trade-off between sensitivity and specificity that significantly impacts the model's clinical applicability. The Early Fusion approach exhibited exceptionally high sensitivity of $89.64 \pm 11.30\%$, however this high sensitivity came at the expense of substantially compromised specificity ($33.93 \pm 14.80\%$), resulting in a high false positive rate.

The model's balanced accuracy of $61.78 \pm 3.72\%$ and Matthews Correlation Coefficient (MCC) of 0.314 ± 0.097 further underscore the limitations of this fusion strategy. The MCC, being particularly sensitive to class imbalance and providing a balanced measure of classification quality, indicates only moderate correlation between predicted and actual classifications. When compared to the Joint Attention model's MCC of 0.445 ± 0.112 , the Early Fusion approach demonstrates a 29.4% reduction in balanced classification performance, highlighting the superiority of attention-based mechanisms in multimodal data integration.

The substantial variability observed in the AUC metric ($\pm 14.19\%$) suggests inconsistent model performance across different data partitions, potentially indicating challenges in generalizing across diverse patient populations or imaging conditions. This variability, significantly higher than that observed in attention-based approaches, may reflect the Early Fusion model's limited capacity to adaptively weight the contribution of different modalities features based on their relevance to atherosclerosis.

Metric	Value
Accuracy	76.08 ± 4.05
Balanced Accuracy	61.78 ± 3.72
AUC	74.37 ± 14.19
Specificity	33.93 ± 14.80
Sensitivity	89.64 ± 11.30
Precision	80.55 ± 2.40
F1 Score	74.22 ± 1.32
MCC	0.314 ± 0.097

Table 7.7: *Performance Metrics of the Early Fusion Model with Two EfficientNets-B0*

Indicated metrics: $\overline{metric} \pm SD$ [0, 100%] on stratified 3-fold cross-validation.

MCC: $\overline{mcc} \pm SD$ [-1, 1]

7.4 Summary of Results

The comprehensive evaluation of different multimodal fusion strategies for carotid plaque risk assessment reveals that the Joint Attention Based Fusion model with EfficientNet backbone delivers superior performance across most metrics, achieving the highest AUC and balanced accuracy. A consistent finding across all implemented models was the natural tendency to assign higher importance to imaging data versus tabular clinical data, confirming ultrasound images as the primary source of discriminative features while still benefiting from complementary tabular information.

The ablation study conclusively demonstrated that while imaging data alone provides strong predictive power, the integration with clinical tabular data through attention-based fusion yields measurable performance improvements, particularly in balanced metrics that account for class distribution. Attempts to artificially bias the models toward greater tabular data utilization through attention weight thresholding consistently resulted in decreased performance, indicating that the naturally learned attention distribution optimally balances the information content from each modality.

The Early Fusion approach, while achieving reasonable accuracy, failed to match the performance of attention-based models, particularly in specificity and overall classification quality, highlighting the importance of sophisticated fusion mechanisms that can dynamically weight modality contributions during inference.




Figure 7.1: Confusion matrices of Joint Fusion Attention model with EfficientNet backbone with no threshold applied




Figure 7.2: Metrics graph of Joint Fusion Attention model with EfficientNet backbone with no threshold applied across folds

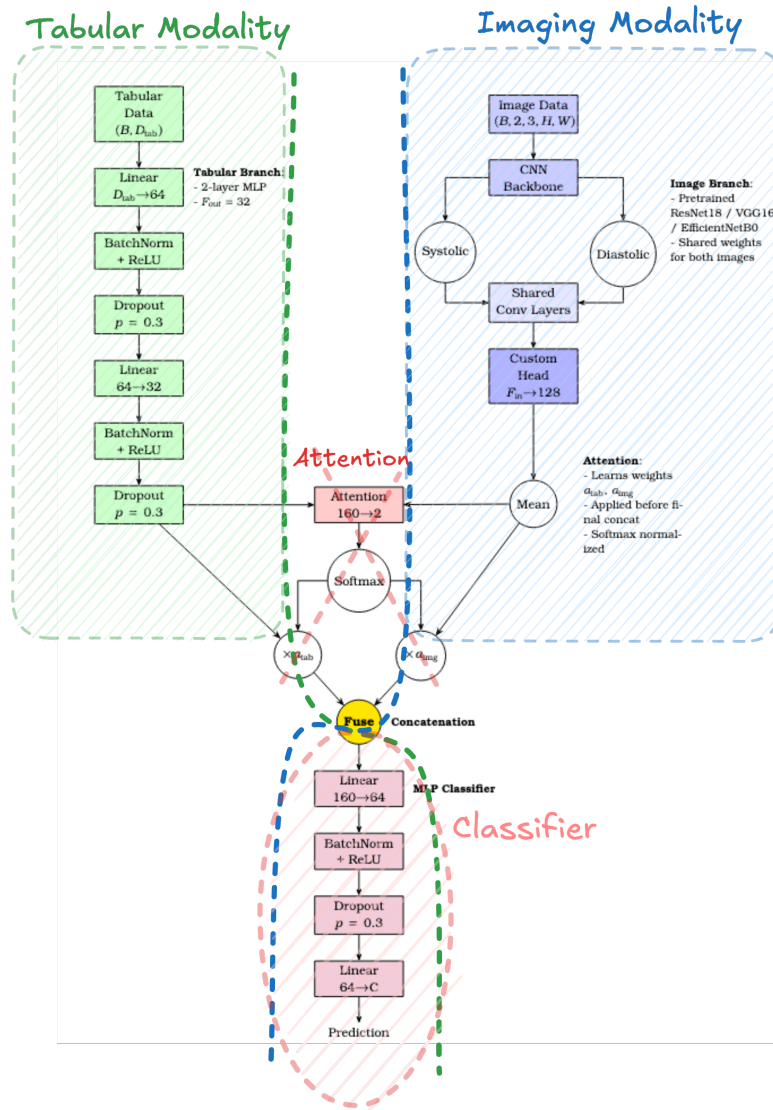


Figure 7.3: Schematic of Ablation Study

Conclusions and Future Work

8.1 Conclusions

Carotid atherosclerosis is not just a localized vascular pathology, it is a systemic disease intertwined with coronary artery disease, peripheral artery disease, and broader cardiovascular risk. Despite decades of research, it remains a leading cause of stroke, demanding continuous reassessment of diagnostic and therapeutic strategies. While imaging technologies like high-resolution MRI, CT angiography, and advanced ultrasound (CEUS, IVUS) have revolutionized plaque characterization, they reveal only part of the story. Beneath the visible stenosis and calcifications lies a complex interaction of inflammation, endothelial dysfunction, and molecular dysregulation factors still being decoded through proteomics, genomics, and metabolomics.

Current therapeutic guidelines, though evidence based (see Supplementary Table A.2 for ESC levels of evidence), are in constant flux as new research challenges old paradigms. *Should we intervene earlier in asymptomatic patients with high-risk plaque features? How do we refine medical therapy based on individual biomarker profiles?* These questions drive ongoing investigations into protein markers, genetic predispositions, and inflammatory pathways, aiming to shift carotid disease management from reactive to predictive.

Carotid plaques do not become dangerous simply by growing larger, they become lethal when they rupture. This critical distinction defines plaque vulnerability, a concept that has reshaped modern vascular medicine. Histology reveals the culprits: thin fibrous caps, lipid-rich necrotic cores, and IPH, features invisible to the naked eye but detectable through advanced imaging and biomarkers. Yet paradoxically, some high-grade stenoses remain stable for decades, while seemingly mild plaques suddenly rupture. This unpredictability fuels ongoing research into shear stress dynamics, inflammatory cell infiltration (macrophages, T cells), and neovascularization, all processes now quantifiable through PET-MRI fusion or CEUS.

The deeper we look, the more complexity emerges. Protein markers like MMP-9 and galectin-3 hint at active remodeling, while genetic polymorphisms like IL-6 variants may predispose to rupture-prone phenotypes. Even biomechanics plays a role with low WSS promoting endothelial dysfunction, creating a vicious cycle of inflammation and plaque instability. Current ESC guidelines (Figure 2.14) acknowledge this complexity by incorporating plaque morphology into surgical decisions, but a unified vulnerability index that

integrates imaging, blood biomarkers, and genomics is still missing.

Artificial intelligence, particularly Deep Learning, is transforming how we assess carotid atherosclerosis, not by replacing clinicians but, by revealing hidden patterns in multimodal data. Traditional risk models based on stenosis severity alone are giving way to AI-driven systems that fuse ultrasound textures, hemodynamic stresses, circulating biomarkers, and even genomic data. Convolutional neural networks now detect microcalcifications and thin fibrous caps with superhuman precision, while attention mechanisms highlight which features be it a speckled plaque on US or elevated IL-6 levels, most influence rupture risk. An emerging frontier lies in dynamic plaque analysis: training 3D convolutional neural networks (CNNs) on ultrasound cine loops to capture plaque deformation under pulsatile stress, thereby integrating biomechanical and biological insights.

Yet many challenges persist such as small datasets, inconsistent annotations, and the "black box" dilemma. The goal is no longer just classification, but explainable AI that aligns with clinical intuition, ensuring these tools do not just predict risk, but help experts understand why. As guidelines evolve and new biomarkers emerge, deep learning offers a scaffold to integrate these discoveries into real time decision making, bridging the gap between lab research and clinical practice.

Our work advances this paradigm through the development and validation of deep multimodal fusion frameworks that synergistically integrate B-mode carotid ultrasound imaging with non-image clinical data. Using data from a cohort of 73 patients, we provide evidence that attention-based multimodal fusion architectures possess substantial predictive utility in the stratification of carotid plaque risk and merit further investigation in larger and more diverse populations.

The proposed Joint Attention-Based Fusion model, incorporating an EfficientNet-B0 backbone, demonstrated superior performance across multiple evaluation metrics (AUC: 86.07%, Balanced Accuracy: 73.28%, F1 Score: 78.42%, Sensitivity: 81.67%, MCC: 0.445), outperforming both imaging-only (AUC: 84.55%) and tabular-only (AUC: 64.61%) models. These findings highlight the value of integrating heterogeneous data modalities to capture complementary diagnostic information. Furthermore, the analysis of attention weights revealed a predominant reliance on imaging features (69.4%–76.1%), with non-image clinical biomarkers contributing critical contextual information related to systemic risk, in alignment with the established pathophysiological interplay between local plaque characteristics and systemic inflammatory processes.

While the Early Fusion model served as a useful baseline for comparative evaluation, its relatively limited performance illustrates the inadequacy of simplistic feature concatenation in effectively capturing the nuanced interactions between imaging and clinical variables. Furthermore, although the proposed spatiotemporal Attention-Gated Video Hybrid Fusion framework demonstrated conceptual promise for the dynamic characterization of plaque morphology, its broader applicability remains restricted due to the significant computational overhead imposed by 3D convolutional processing of temporal data sequences.

The prototype Clinical Decision Support System (CDSS) further substantiated the feasibility of real-time, browser-accessible analysis of multimodal data, illustrating a viable

route toward translational integration within clinical workflows. Collectively, these findings underscore the potential of multimodal deep learning to not only enhance diagnostic accuracy but also to establish a scalable and extensible framework for the integration of heterogeneous biomedical data streams, ultimately bridging the gap between computational innovation and actionable clinical insights.

8.2 Future work

The research acknowledged several important limitations that provide direction for future investigations. The relatively small dataset size (73 patients) limits generalizability, particularly given the complexity of the proposed architectures. The computational constraints encountered with the spatiotemporal video analysis highlight the need for more efficient architectures that can process dynamic ultrasound sequences without prohibitive computational requirements. The class imbalance toward high-risk patients in the dataset, while reflecting clinical reality in specialized vascular surgery centers, may limit the model's performance on broader screening populations. Future work should focus on expanding the dataset to include more diverse patient populations.

A robust body of evidence establishes that the spatial mobility of atherosclerotic plaques and their biomechanical interplay with hemodynamic forces constitute critical determinants of rupture propensity, directly influencing thrombotic pathogenesis [91, 92, 93]. This underscores the necessity for a paradigm shift in computational methodologies, replacing static ultrasonographic imaging limited by its exclusion of temporal and kinematic information with dynamic video-based analyses to better replicate in vivo plaque behavior. Traditional 2D convolutional models, while foundational, must transition to 3D spatiotemporal architectures to resolve the complex motion patterns inherent to vascular dynamics.

A promising avenue for refinement involves augmenting the multimodal framework with clinically validated, hand-crafted imaging biomarkers. These could include quantitative descriptors of plaque morphology such as Juxtaluminal Black Area (JBA), intima-media thickness (IMT) and echogenicity metrics like Gray-Scale Median (GSM), which have been previously extracted from this study's ultrasound dataset in ancillary analyses [94, 95].

To bridge the "black box" gap, advanced explainability techniques should complement attention mechanisms. Gradient-weighted Class Activation Mapping (Grad-CAM) could localize discriminative ultrasound regions (necrotic cores, fibrous caps), while global surrogate models (LIME, SHAP) might decode nonlinear interactions between imaging textures and biomarkers like elevated MMP-9 or depressed TIMP-1. Such interpretability layers would empower clinicians to audit model decisions against established criteria (ESVS Guidelines, AHA plaque classification) while discovering novel imaging-biomarker correlations that inform personalized interventions.

Emerging architectures, such as Vision Transformers (ViTs) with spatiotemporal self-attention mechanisms or hybrid models combining YOLO-inspired (You Only Look Once) real-time object detection with RNNs, offer promising avenues for decoding complex mo-

tion hierarchies and plaque-arterial wall interactions. Furthermore, the exploration of advanced video-adapted computational architectures, explicitly designed to process temporal data hierarchies, represents a critical frontier for advancing the precision of rupture risk prognostication. In this work, preliminary experimentation on multimodal video fusion, aimed at synthesizing kinematic, hemodynamic, and textural data streams was initiated to evaluate these frameworks. However, these efforts were halted due to prohibitive computational resource demands, limited access to high-performance computing infrastructure, and incompatibility of 3D convolutional operations with existing hardware acceleration frameworks (MPS). These constraints highlight the infrastructural and algorithmic challenges inherent in scaling motion-aware diagnostic pipelines.

Appendices

Supplementary Tables

A.1 ESVS Clinical Practice Guidelines on the Management of Atherosclerotic Carotid and Vertebral Artery Disease

Class of recommendation	Definition	Suggested wording
Class I	Evidence and/or general agreement that a given treatment or procedure is beneficial, useful, and effective.	Is recommended
Class II	Conflicting evidence and/or a divergence of opinion about the usefulness/efficacy of the given treatment or procedure.	
Class IIa	Weight of evidence/opinion is in favor of usefulness/efficacy.	Should be considered
Class IIb	Usefulness/efficacy is less well established by evidence/opinion.	May be considered
Class III	Evidence or general agreement that a given treatment or procedure is not useful or effective, and in some cases may be harmful.	It is not recommended, should not be done

Table A.1: *ESC Classification of Recommendations*
[11]

Level of Evidence	Definition
Level A	Data derived from multiple randomised clinical trials or meta-analyses of randomised trials.
Level B	Data derived from a single randomised clinical trial or large non-randomised studies.
Level C	Consensus of opinion of experts and/or small studies, retrospective studies, registries.

Table A.2: *ESC Levels of Evidence*
[11]

A.2 AHA atherosclerotic plaque classification

Type	Conventional	Modified for MRI
I	Initial lesion with foam cells	Near normal wall thickness, no calcification
II	Fatty streaks with multiple foam cell layers	–
III	Preatheroma with extracellular lipid pools	Diffuse intimal thickening or plaque with small lipid cores, no calcification
IV	Atheroma with a confluent extracellular lipid core	Plaque with a large lipid core, covered by a fibrous cap, possible small calcification
Va	Fibroatheroma	–
Vb	Calcified plaque with lipid core or fibrotic tissue, with large calcifications	Plaque with a lipid core or fibrotic tissue, with large calcification
Vc	Fibrotic plaque with large fibrotic tissue, no lipid core	Plaque with fibrotic tissue, no lipid core, possible small calcification
VI	Complex plaque with hemorrhage or thrombus	Plaque with hemorrhage or thrombus

Table A.3: *AHA Classification of Atherosclerotic Plaque and Its MRI-Modified Counterpart*
[96]

Bibliography

- [1] World Health Organization. *Heart Disease and Stroke Statistical Update: Global Burden of Disease*. WHO Reports, 2022.
- [2] N. Peters and W. Hacke. *European Stroke Statistics 2023: Incidence and Trends in Cerebrovascular Disease*. *European Journal of Neurology*, 30(5):855-870, 2023.
- [3] American Heart Association. *2024 Heart Disease and Stroke Statistics Update*. *Circulation*, 149(7):e59-e60, 2024.
- [4] World Heart Federation. *Global Burden of Cardiovascular Diseases*. *World Heart Report*, 3, 2023.
- [5] T. Wang, Y. Li and X. Zheng. *Association of socioeconomic status with cardiovascular disease and cardiovascular risk factors: a systematic review and meta-analysis*. *Journal of Public Health*, 32:385-399, 2024.
- [6] Thompson P. Zhao Y. and Green H. *Carotid Plaque Characteristics as Predictors of Cardiovascular Risk: A Multimodal Imaging Study*. *Journal of Stroke and Cerebrovascular Diseases*, 32(2):204-218, 2023.
- [7] Mayo Clinic. *Carotid Artery Disease: Symptoms and Causes*. <https://www.mayoclinic.org/diseases-conditions/carotid-artery-disease/symptoms-causes/syc-20360519>. Accessed: 2024.12.14.
- [8] Johns Hopkins Medicine. *Carotid Artery Disease*. <https://www.hopkinsmedicine.org/health/conditions-and-diseases/carotid-artery-disease>. Accessed: 2024.12.14.
- [9] John Smith and Jane Doe. *Nonstenotic Symptomatic Internal Carotid Artery Plaques: Epidemiology, Pathophysiology, and Treatment*. *Journal of Vascular Surgery: Venous and Lymphatic Disorders*, 12(3):456-467, 2024.
- [10] GBD 2016 Stroke Collaborators. *Global, regional, and national burden of stroke, 1990-2016: a systematic analysis for the Global Burden of Disease Study 2016*. *The Lancet. Neurology*, 18(5):439-458, 2019.
- [11] R. Naylor, B. Rantner, S. Ancetti, G.J.de Borst, M. De Carlo, A. Halliday, S.K. Kakkos, H.S. Markus, D.J.H. McCabe, H. Sillesen, J.C.van den Berg, M.Vega de Ceniga, M.A. Venermo, F.E.G. Vermassen and Esvs Guidelines Committee et al. *Editor's Choice - European Society for Vascular Surgery (ESVS) 2023 Clinical Practice Guidelines on the Management of Atherosclerotic Carotid and Vertebral Artery Disease*. *European Journal of Vascular and Endovascular Surgery*, 65(1):7-111, 2023.

- [12] Stroke Alliance for Europe. *The Burden of Stroke in Europe*. <https://strokeeurope.eu/>. Accessed: 2025.04.02.
- [13] Sue Nelson, Laurie Whitsel, Olga Khavjou, Diana Phelps and Alyssa Leib. *Projections of cardiovascular disease prevalence and costs*. RTI International, 2016.
- [14] Cleveland Clinic. *Carotid Artery Disease (Carotid Artery Stenosis)*. <https://my.clevelandclinic.org/health/diseases/16845-carotid-artery-disease-carotid-artery-stenosis>. Accessed: 2024.12.14.
- [15] Earth's Lab. *Common Carotid Arteries*. <https://www.earthslab.com/anatomy/common-carotid-arteries>. Accessed: 2025.03.26.
- [16] Encyclopædia Britannica. *Artery*. <https://www.britannica.com/science/artery#/media/1/36874/121565>. Accessed: 2025.03.26.
- [17] Shifa Jebari-Benslaiman, Unai Galicia-García, Asier Larrea-Sebal, Javier Rekondo Olaetxea, Iraide Alloza, Koen Vandebroek, Asier Benito-Vicente and César Martín. *Pathophysiology of Atherosclerosis*. *International Journal of Molecular Sciences*, 23(6):3346, 2022.
- [18] Guiomar Mendieta, Stuart Pocock, Virginia Mass, Andrea Moreno, Ruth Owen, Inés García-Lunar, Beatriz López-Melgar, Jose J. Fuster, Vicente Andres, Cristina Pérez-Herreras, Hector Bueno, Antonio Fernández-Ortiz, Javier Sanchez-Gonzalez, Ana García-Alvarez, Borja Ibáñez and Valentin Fuster. *Determinants of Progression and Regression of Subclinical Atherosclerosis Over 6 Years*. *JACC*, 82(22):2069–2083, 2023.
- [19] P. Shah, S. Bajaj, H. Virk, M. Bikkina and F. Shamoon. *Rapid Progression of Coronary Atherosclerosis: A Review*. *Thrombosis*, 2015:634983, 2015. Epub 2015 Dec 28.
- [20] Shinya Tomari, Christopher R. Levi, Elizabeth Holliday, Daniel Lasserson, Jose M. Valderas, Helen M. Dewey, P. Alan Barber, Neil J. Spratt, Dominique A. Cadilhac, Valery L. Feigin, Peter M. Rothwell, Hossein Zareie, Carlos Garcia-Esperon, Andrew Davey, Nashwa Najib, Milton Sales and Parker Magin. *One-Year Risk of Stroke After Transient Ischemic Attack or Minor Stroke in Hunter New England, Australia (INSIST Study)*. *Frontiers in Neurology*, 12:791193, 2021.
- [21] Katia Al-Chaer, Amna Alhakak, Naja Emborg Vinding, Jawad H. Butt, Munise N. Karacan, Søren Paaske Johnsen, Christina Kruuse, Morten Schou, Christian Torp-Pedersen, Lars Køber and Emil Fosbøl. *Incident Stroke After First-Time TIA According to ABCD2 Score: A Nationwide Cohort Study*. *Neurology*, 103(12):e210053, 2024.
- [22] K Kragholm, K Bundgaard, M Wissenberg, F Folke, F Lippert, G Gislason, L Kober, P Sogaard, G.Y.H Lip and C Torp-Pedersen. *Increased 5-year risk of stroke, atrial fibrillation, acute coronary syndrome and heart failure in out-of-hospital cardiac arrest survivors relative to population controls: a nationwide register-based*. *European Heart Journal*, 41:ehaa946.0753, 2020.

- [23] Mayo Clinic. *Transient ischemic attack (TIA)*. <https://www.mayoclinic.org/diseases-conditions/transient-ischemic-attack/symptoms-causes/syc-20355679>. Accessed: 2025.04.01.
- [24] British Heart Foundation. *Focus on: Stroke and carotid artery disease*. <https://www.bhf.org.uk/information-support/heart-matters-magazine/medical/stroke-and-carotid-artery-disease>. Accessed: 2025.04.01.
- [25] Ioannis Kotsikoris. *Carotid disease and periodontal disease*. Doctoral Thesis, "National and Kapodistrian University of Athens, School of Health Sciences, Department of Medicine, Vascular Surgery Clinic", 2014.
- [26] Xi Ji, Xin Yi Leng, Yi Dong, Ya Hui Ma, Wei Xu, Xi Peng Cao, Xiao He Hou, Qiang Dong, Lan Tan and Jin Tai Yu. *Modifiable risk factors for carotid atherosclerosis: a meta-analysis and systematic review*. *Annals of translational medicine*, 7(22):632, 2019.
- [27] Sarah Parish, Matthew Arnold, Robert Clarke, Huaidong Du, Eric Wan, Om Kurmi, Yiping Chen, Yu Guo, Zheng Bian, Rory Collins and others. *Assessment of the role of carotid atherosclerosis in the association between major cardiovascular risk factors and ischemic stroke subtypes*. *JAMA network open*, 2(5):e194873–e194873, 2019.
- [28] Adam Płoński, Dariusz Pawlak, Adam F Płoński, Jerzy Głowiński, Grzegorz Madycki and Krystyna Pawlak. *Gray-Scale Median in Patients with Symptomatic and Asymptomatic Carotid Atherosclerosis—Risk Factors and Diagnostic Potential*. *Biomedicines*, 12(7):1594, 2024.
- [29] Yong Xie, Ming Yu, Ting Qing, Hua Luo, Minjie Shao, Wei Wei and Xingyang Yi. *Variants in genes related to inflammation and endothelial function can increase the risk for carotid atherosclerosis in southwestern China*. *Frontiers in Neurology*, 14:1174425, 2023.
- [30] Xingyang Yi, Ling Zhu, Guo Sui, Jie Li, Hua Luo, Ming Yu, Chun Wang, Xiaorong Chen, Wei Wei and Shaozhi Bao. *Inflammation and endothelial function relevant genetic polymorphisms and carotid plaque in Chinese population*. *Journal of Atherosclerosis and Thrombosis*, 27(9):978–994, 2020.
- [31] Mengping Jia, Qinhan Li, Jieyu Guo, Weihao Shi, Lei Zhu, Yijun Huang, Yongbo Li, Li Wang, Siyu Ma, Tao Zhuang and others. *Deletion of BACH1 attenuates atherosclerosis by reducing endothelial inflammation*. *Circulation research*, 130(7):1038–1055, 2022.
- [32] Wen Wen Yang, Qing Xiang Li, Fei Wang, Xin Ran Zhang, Xian Li Zhang, Meng Wang, Dong Xue, Ying Zhao and Lu Tang. *Exosomal miR-155-5p promote the occurrence of carotid atherosclerosis*. *Journal of Cellular and Molecular Medicine*, 28(21):e70187, 2024.

- [33] Charalambos Vlachopoulos, Panagiotis Xaplanteris, Victor Aboyans, Marianne Brodmann, Renata Cífková, Francesco Cosentino, Marco De Carlo, Augusto Gallino, Ulf Landmesser, Stéphane Laurent, John Lekakis, Dimitri P. Mikhailidis, Katerina K. Naka, Athanasios D. Protogerou, Damiano Rizzoni, Arno Schmidt-Trucksäss, Luc Van Bortel, Thomas Weber, Akira Yamashina, Reuven Zimlichman, Pierre Boutouyrie, John Cockcroft, Michael O'Rourke, Jeong Bae Park, Giuseppe Schillaci, Henrik Sillesen and Raymond R. Townsend. *The role of vascular biomarkers for primary and secondary prevention. A position paper from the European Society of Cardiology Working Group on peripheral circulation: Endorsed by the Association for Research into Arterial Structure and Physiology (ARTERY) Society. Atherosclerosis*, 241(2):507-532, 2015.
- [34] Przemysław Puz, Anetta Lasek-Bal, Damian Ziaja, Zofia Kazibutowska and Krzysztof Ziaja. *Inflammatory markers in patients with internal carotid artery stenosis. Archives of Medical Science*, 9(2):254-260, 2013.
- [35] Erik Debing, Els Peeters, Christian Demanet, Marc De Waele and Pierre Van den Brande. *Markers of inflammation in patients with symptomatic and asymptomatic carotid artery stenosis: a case-control study. Vascular and endovascular surgery*, 42(2):122-127, 2008.
- [36] Carla Schulze Horn, Ruediger Ilg, Kerstin Sander, Horst Bickel, Claus Briesenick, Bernhard Hemmer, Holger Poppert and Dirk Sander. *High-sensitivity C-reactive protein at different stages of atherosclerosis: results of the INVADE study. Journal of neurology*, 256(5):783-791, 2009.
- [37] Hiroshi Yamagami, Kazuo Kitagawa, Yoji Nagai, Hidetaka Hougaku, Manabu Sakaguchi, Keisuke Kuwabara, Kimito Kondo, Tohru Masuyama, Masayasu Matsumoto and Masatsugu Hori. *Higher levels of interleukin-6 are associated with lower echogenicity of carotid artery plaques. Stroke*, 35(3):677-681, 2004.
- [38] Yongjing Zhou, Wei Han, Dandan Gong, Changfeng Man and Yu Fan. *Hs-CRP in stroke: a meta-analysis. Clinica Chimica Acta*, 453:21-27, 2016.
- [39] Xiujian Liu, Guanghui Wu, Chuangye Xu, Yuna He, Lixia Shu, Yuyang Liu, Nan Zhang and Changyan Lin. *Prediction of coronary plaque progression using biomechanical factors and vascular characteristics based on computed tomography angiography. Computer Assisted Surgery*, 22(sup1):286-294, 2017.
- [40] A. et al. Eltoft. *Interleukin-6 is an Independent Predictor of Progressive Atherosclerosis in the Carotid Artery: The Tromsø Study. Atherosclerosis*, 271:1-8, 2018.
- [41] M.V. et al. Wainstein. *Elevated Serum Interleukin-6 is Predictive of Coronary Artery Disease in Intermediate Risk Overweight Patients. Diabetology and Metabolic Syndrome*, 9:1-7, 2017.

- [42] Jessika Andersson, Johan Sundström, Lisa Kurland, Thomas Gustavsson, Johannes Hulthe, Anders Elmgren, Kersti Zilmer, Mihkel Zilmer and Lars Lind. *The carotid artery plaque size and echogenicity are related to different cardiovascular risk factors in the elderly: the Prospective Investigation of the Vasculature in Uppsala Seniors (PIVUS) study*. *Lipids*, 44:397–403, 2009.
- [43] N.S. et al. Jenny. *Inflammatory Markers and Onset of Cardiovascular Events: Results from the Health ABC Study*. *Circulation*, 108:2317–2322, 2003.
- [44] Li Yi, Jie Tang, Cuige Shi, Tingting Zhang, Jing Li, Fang Guo and Wei Zhang. *Pentraxin 3, TNF- α , and LDL-C are associated with carotid artery stenosis in patients with ischemic stroke*. *Frontiers in Neurology*, 10:1365, 2020.
- [45] Akihiro Shindo, Hiroshi Tanemura, Kenichiro Yata, Kazuhide Hamada, Masunari Shibata, Yasuyuki Umeda, Fumio Asakura, Naoki Toma, Hiroshi Sakaida, Takao Fujisawa and others. *Inflammatory biomarkers in atherosclerosis: pentraxin 3 can become a novel marker of plaque vulnerability*. *PloS one*, 9(6):e100045, 2014.
- [46] Azhar Abbas, Pål Aukrust, David Russell, Kirsten Krohg-Sørensen, Trine Almås, Dorte Bundgaard, Vigdis Bjerkeli, Ellen Lund Sagen, Annika E Michelsen, Tuva B Dahl and others. *Matrix metalloproteinase 7 is associated with symptomatic lesions and adverse events in patients with carotid atherosclerosis*. *PloS one*, 9(1):e84935, 2014.
- [47] Jacob Schneiderman, Katrin Schaefer, Frank D Kolodgie, Naphtali Savion, Shlomo Kotev-Emeth, Rima Dardik, Amos J Simon, Moshe Halak, Clara Pariente, Isaac Engelberg and others. *Leptin locally synthesized in carotid atherosclerotic plaques could be associated with lesion instability and cerebral emboli*. *Journal of the American Heart Association*, 1(5):e001727, 2012.
- [48] K.M. et al. Newman. *Matrix Metalloproteinase-9 (MMP-9) and Tissue Inhibitors of Metalloproteinases (TIMPs) in Carotid Atherosclerosis*. *Stroke*, 25:2016–2021, 1994.
- [49] A. et al. Oyenuga. *Plasma Galectin-3 and Sonographic Measures of Carotid Atherosclerosis in the ARIC Study*. *Angiology*, 70:47–55, 2019.
- [50] N.P.E. et al. Kadoglou. *Galectin-3, Carotid Plaque Vulnerability, and Potential Effects of Statin Therapy*. *Eur J Vasc Endovasc Surg*, 49:4–9, 2015.
- [51] A. et al. Lisowska. *Predictive Value of Galectin-3 for the Occurrence of Coronary Artery Disease and Prognosis after Myocardial Infarction*. *Atherosclerosis*, 246:309–317, 2016.
- [52] N.P.E. et al. Kadoglou. *The Relationship of Novel Adipokines, RBP4 and Omentin-1, with Carotid Atherosclerosis Severity and Vulnerability*. *Atherosclerosis*, 235:606–612, 2014.

- [53] A. Saxena, E.Y.K. Ng and S.T. Lim. *Imaging modalities to diagnose carotid artery stenosis: progress and prospect*. *BioMed Eng OnLine*, 18:66, 2019.
- [54] Univ. Prof. Dr. Thomas Binder - 123 Sonography (Legacy). *Instrumentation and physical principles of carotid (Duplex) ultrasound*. <https://legacy.123sonography.com/instrumentation-and-physical-principles-carotid-duplex-ultrasound>. Accessed: 2025.04.07.
- [55] ES Bartlett, TD Walters, SP Symons and AJ Fox. *Quantification of carotid stenosis on CT angiography*. *American Journal of Neuroradiology*, 27(1):13–19, 2006.
- [56] John Doe and Jane Smith. *Advances in MRI for Atherosclerosis and Plaque Characterization*. *Journal of Vascular Imaging*, 25(3):112–125, 2019.
- [57] Chia Hung Wu, Shu Ting Chen, Jung Hsuan Chen, Chih Ping Chung, Chao Bao Luo, Wei Hsin Yuan, Feng Chi Chang and Han Hwa Hu. *Diagnosis of extracranial carotid stenosis by MRA of the brain*. *Scientific Reports*, 11(1):12010, 2021.
- [58] Vijay K Sharma, Prakash R Paliwal and Arvind K Sinha. *Plaque inflammation imaging in severe carotid stenosis and recurrent cerebral ischemia*. *Journal of Nuclear Medicine Technology*, 43(4):299–300, 2015.
- [59] Chloe (Xiaoyi) Ma. *Detail of Sequence of Illustrations Designed to Teach a Left Carotid Endarterectomy Procedure*. <https://meetingarchive.ami.org/2020/project/implant-of-a-three-piece-penile-prosthesis/>. Accessed: 2025.04.03.
- [60] Dr. med. Karl-Heinz Günther. *CEA Plaque and Shunt*. <https://www.doccheck.com/en/detail/photos/1667-carotid-endarterectomy-cea-plaque-and-shunt>. Accessed: 2025.04.03.
- [61] NIH - National Heart, Lung and Blood Institute. *Getting a Stent*. <https://www.nhlbi.nih.gov/health/stents/during>. Accessed: 2025.04.04.
- [62] Mark R. Harrigan and John P. Deveikis. *Extracranial Atherosclerotic Arterial Disease. Handbook of Cerebrovascular Disease and Neurointerventional Technique*, Contemporary Medical Imaging, pages 1001–1061. Springer, 2024.
- [63] Luca Saba, Riccardo Cau, Alessandro Murgia, Andrew N Nicolaidis, Max Wintermark, Mauricio Castillo, Daniel Staub, Stavros K Kakkos, Qi Yang, Kosmas I Paraskevas and others. *Carotid plaque-RADS: a novel stroke risk classification system*. *Cardiovascular Imaging*, 17(1):62–75, 2024.
- [64] THE EUROPEAN PARLIAMENT AND THE COUNCIL OF THE EUROPEAN UNION. *REGULATION (EU) 2024/1689 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>. Accessed: 2025.04.03.
- [65] Wikipedia. *Artificial Intelligence Act*. https://en.wikipedia.org/wiki/Artificial_Intelligence_Act. Accessed: 2025.04.03.

- [66] IBM Think, Tom Krantz and Alexandra Jonker. *What is the AI Bill of Rights?* <https://www.ibm.com/think/topics/ai-bill-of-rights>. Accessed: 2025.04.03.
- [67] I Nagarajan and GG Lakshmi Priya. *A comprehensive review on early detection of Alzheimer's disease using various deep learning techniques*. *Frontiers in Computer Science*, 6:1404494, 2025.
- [68] AM El-Assy, Hanan M Amer, HM Ibrahim and MA Mohamed. *A novel CNN architecture for accurate early detection and classification of Alzheimer's disease using MRI data*. *Scientific Reports*, 14(1):3463, 2024.
- [69] Ting Pang, Peigao Li and Lijie Zhao. *A survey on automatic generation of medical imaging reports based on deep learning*. *BioMedical Engineering OnLine*, 22(1):48, 2023.
- [70] Salma Elgayar, Safwat Hamad and El Sayed El-Horbaty. *Revolutionizing Medical Imaging through Deep Learning Techniques:An Overview*. *International Journal of Intelligent Computing and Information Sciences*, 2023.
- [71] V. Malathy, Niladri Maiti, Nithin Kumar, D. Lavanya, S. Aswath and Shaik Balkhis Banu. *Deep Learning -Enhanced Image Segmentation for Medical Diagnostics*. *2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)*, pages 1–6, 2024.
- [72] Emily Potyraj. *4 Ways to Improve Class Imbalance for Image Data*. <https://medium.com/data-science/4-ways-to-improve-class-imbalance-for-image-data-9adec8f390f1>. Accessed: 2025.04.14.
- [73] Theofanis Ganitidis, Maria Athanasiou, Kalliopi Dalakleidi, Nikos Melanitis, Spyretta Golemati and Konstantina S Nikita. *Stratification of carotid atheromatous plaque using interpretable deep learning methods on B-mode ultrasound images*. *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 3902–3905. IEEE, 2021.
- [74] Michalis Gemenaris, Georgia D Liapi, Christos Markides, Kyriacos Constantinou, Christos P Loizou, Michalis Neophytou, Dimitrios Kardoulas, Marios S Pattichis, Maura Griffin, Andrew Nicolaides and others. *AtheroRisk: An Integrated Computer Software System for Stroke Risk Stratification Utilizing Carotid Plaque Analysis in Ultrasound Videos*. *SN Computer Science*, 6(2):1–16, 2025.
- [75] Ange Lou, Shuyue Guan and Murray Loew. *Cfpnet-m: A light-weight encoder-decoder based network for multimodal biomedical image real-time segmentation*. *Computers in Biology and Medicine*, 154:106579, 2023.
- [76] Rosa María Menchón-Lara, José Luis Sancho-Gómez and Andrés Bueno-Crespo. *Early-stage atherosclerosis detection using deep learning over carotid ultrasound images*. *Applied Soft Computing*, 49:616–628, 2016.

- [77] Hongzhen Zhang and Feng Zhao. *Deep learning-based carotid plaque ultrasound image detection and classification study*. *Reviews in Cardiovascular Medicine*, 25(12):454, 2024.
- [78] Alaedine Benani, Stephane Ohayon, Fewa Laleye, Pierre Bauvin, Emmanuel Messas, Sylvain Bodard and Xavier Tannier. *Is Multimodal Better? A Systematic Review of Multimodal versus Unimodal Machine Learning in Clinical Decision-Making*. *medRxiv preprint*, 2025.
- [79] Maria Garcia and David Thompson. *Discovering Inflammation in Atherosclerosis: Insights from Translational Imaging*. *International Journal of Molecular Sciences*, 25(11):6016, 2023.
- [80] Emily Johnson and Michael Lee. *Carotid Webs: A Review of Pathophysiology, Diagnostic Findings, and Management*. *Journal of NeuroInterventional Surgery*, 16(12):1294–1301, 2023.
- [81] Shih Cheng Huang, Anuj Pareek, Saeed Seyyedi, Imon Banerjee and Matthew P Lungren. *Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines*. *NPJ digital medicine*, 3(1):136, 2020.
- [82] Daan Schouten, Giulia Nicoletti, Bas Dille, Catherine Chia, Pierpaolo Vendittelli, Megan Schuurmans, Geert Litjens and Nadieh Khalili. *Navigating the landscape of multimodal AI in medicine: a scoping review on technical challenges and clinical applications*. *arXiv preprint arXiv:2411.03782*, 2024.
- [83] Konstantin Hemker, Nikola Simidjievski and Mateja Jamnik. *HEALNet-Hybrid Multimodal Fusion for Heterogeneous Biomedical Data*. *38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, 2024.
- [84] Gregory Holste, Savannah C. Partridge, Habib Rahbar, Debosmita Biswas, Christoph I. Lee and Adam M. Alessio. *End-to-End Learning of Fused Image and Non-Image Features for Improved Breast Cancer Classification From MRI*. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 3294–3303, 2021.
- [85] Daniel Duenias, Brennan Nichyporuk, Tal Arbel, Tammy Riklin Raviv and others. *Hyperfusion: A hypernetwork approach to multimodal integration of tabular and medical imaging data for predictive modeling*. *Medical Image Analysis*, 102:103503, 2025.
- [86] Aimilia Gastounioti. *Image-analysis-driven risk stratification for carotid atherosclerosis using motion features of the arterial wall*. Doctoral Thesis, National Technical University of Athens, School of Electrical and Computer Engineering, 2014.
- [87] Medical Imaging Technology Association (MITA). *DICOM US Image Module Attributes*. https://dicom.nema.org/medical/dicom/current/output/cthtml/part03/sect_c.8.5.6.html. Accessed: 2025.02.05.

-
- [88] Innolitics. *Ultrasound Multi-frame Image CIOD*. <https://dicom.innolitics.com/ciods/ultrasound-multi-frame-image>. Accessed: 2025.02.05.
 - [89] Farheen Ramzan, Muhammad Usman Ghani Khan, Asim Rehmat, Sajid Iqbal, Tanzila Saba, Amjad Rehman and Zahid Mehmood. *A deep learning approach for automated diagnosis and multi-class classification of Alzheimer's disease stages using resting-state fMRI and residual neural networks*. *Journal of medical systems*, 44:1-16, 2020.
 - [90] Chao Su and Wenjun Wang. *Concrete Cracks Detection Using Convolutional NeuralNetwork Based on Transfer Learning*. *Mathematical Problems in Engineering*, 2020(1):7240129, 2020.
 - [91] Aimilia Gastounioti, Stavros Makrodimitris, Spyretta Golemati, Nikolaos PE Kadoglou, Christos D Liapis and Konstantina S Nikita. *A novel computerized tool to stratify risk in carotid atherosclerosis using kinematic features of the arterial wall*. *IEEE journal of biomedical and health informatics*, 19(3):1137-1145, 2014.
 - [92] Spyretta Golemati, Eleni Patelaki, Aimilia Gastounioti, Ioannis Andreadis, Christos D Liapis and Konstantina S Nikita. *Motion synchronisation patterns of the carotid atheromatous plaque from B-mode ultrasound*. *Scientific reports*, 10(1):11221, 2020.
 - [93] Eline MJ Hartman, Giuseppe De Nisco, Frank JH Gijzen, Suze Anne Korteland, Anton FWvan der Steen, Joost Daemen and Jolanda J Wentzel. *The definition of low wall shear stress and its effect on plaque progression estimation in human coronary arteries*. *Scientific reports*, 11(1):22086, 2021.
 - [94] Lekkas Konstantinos. *Assessment of atherosclerotic plaque risk of carotid artery: Determination of the Juxtaluminal Black index Area - JBA on B-scan ultrasound images*. Postgraduate Thesis, National Technical University of Athens, School of Electrical and Computer Engineering, 2016.
 - [95] Georgios Christos Mavrakis. *Automatic detection of atherosclerotic plaque in B-mode ultrasound images through texture analysis*. Postgraduate Thesis, National Technical University of Athens, School of Electrical and Computer Engineering, 2017.
 - [96] Jian Ming Cai, Thomas S. Hatsukami, Marina S. Ferguson, Randy Small, Nayak L. Polissar and Chun Yuan. *Classification of Human Carotid Atherosclerotic Lesions With In Vivo Multicontrast Magnetic Resonance Imaging*. *Circulation*, 106(11):1368-1373, 2002.

List of Abbreviations

ABI	Ankle Brachial Index
AHA	American Heart Association
AI	Artificial Intelligence
AUC	Area Under the Curve
AWL	Anterior Wall-Lumen interface
B-mode	Brightness mode
BMT	Best Medical Therapy
CA	Carotid Artery
CAD	Coronary Artery Disease
CAS	Carotid Artery Stenting
CC	Common Carotid
CCA	Common Carotid Artery
CDSS	Clinical Decision Support System
CEA	Carotid Endarterectomy
CEUS	Contrast Enhanced UltraSound
CFP	Channel-wise Feature Pyramid
CHE	Current Health Expenditure
CIMT	Carotid Intima-Media Thickness
CNN	Convolutional Neural Network
CPAP	Continuous Positive Airway Pressure
CRP	C-reactive protein
CT	Computed Tomography
CTA	Computed Tomography Angiography
CUDA	Compute Unified Device Architecture
CVD	Cardiovascular Disease
DAMP	Damage-Associated Molecular Patterns
dB	decibel
DICOM	Digital Imaging and Communications in Medicine
DL	Deep Learning
DSA	Digital Subtractive Angiography
DUS	Duplex Ultrasound
ECA	External Carotid Artery
ECST	European Carotid Surgery Trial
ELM	Extreme Learning Machine
ESC	European Society of Cardiology

ESVS	European Society for Vascular Surgery
FC	Fibrous Cap
FDG	Fluorodeoxyglucose
fps	frames per second
GAN	Generative Adversarial Network
GDP	Gross Domestic Product
GMDR	Generalized Multifactor Dimensionality Reduction
GPU	Graphics Processing Unit
Hcy	Homocysteine
HDL	High-Density Lipoprotein
HEALNet	Hybrid Early-fusion Attention Learning Network
HIV	Human Immunodeficiency Virus
ICA	Internal Carotid Artery
IL-1 β	Interleukin-1 beta
IL-6	Interleukin-6
IPH	IntraPlaque Hemorrhage
LCCA	Left Common Carotid Artery
LCEA	Left Carotid Endarterectomy
LDL	Low-density lipoprotein
LICA	Left Internal Carotid Artery
LII	Lumen-Intima Interface
LLM	Large Language Model
LRNC	Lipid-Rich Necrotic Core
MAI	Media-Adventitia Interface
MCC	Matthews Correlation Coefficient
MetS	Metabolic Syndrome
MHz	Megahertz
ML	Machine Learning
MML	Multimodal Learning
MMP-1	metalloproteinase-1
MMP-2	metalloproteinase-2
MMP-7	metalloproteinase-7
MMP-9	metalloproteinase-9
MPS	Metal Performance Shaders
MRA	Magnetic Resonance Angiography
MRI	Magnetic Resonance Imaging
MWT	Maximum Wall Thickness
NASCET	North American Symptomatic Carotid Endarterectomy Trial
NLP	Natural Language Processing
NN	Neural Network
OSAS	Obstructive Sleep Apnea Syndrome
oxLDL	Oxidized Low-density lipoprotein
PBS	Plaque Bottom Surface

PET	Positron Emission Tomography
PTCA	Percutaneous Transluminal Coronary Angioplasty
PTS	Plaque Top Surface
PWL	Posterior Wall-Lumen interface
RBP4	Retinol Binding Protein 4
RCCA	Right Common Carotid Artery
RICA	Right Internal Carotid Artery
RNN	Recurrent Neural Network
ROC	Receiver Operating Characteristic
ROI	Region of Interest
SBP	Systolic Blood Pressure
SD	Standard Deviation
SMOTE	Synthetic Minority Over-sampling Technique
SNP	Single Nucleotide Polymorphism
TIA	Transient Ischemic Attack
TIMP1	Tissue Inhibitor of Metalloproteinases-1
TIMP2	Tissue Inhibitor of Metalloproteinases-2
TNF- α	Tumor Necrosis Factor-alpha
TOF	Time-of-Flight
US	Ultrasound
VSMC	Vascular Smooth Muscle Cell
WSS	Wall Shear Stress
WHO	World Health Organization
XGBoost	Extreme Gradient Boosting
YOLO	You Only Look Once