UNIVERSIDADE FEDERAL DE UBERLÂNDIA
FACULDADE DE ENGENHARIA ELÉTRICA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

DEEP LEARNING-BASED DETECTION OF CAROTID ARTERY ATHEROMAS IN PANORAMIC RADIOGRAPHS

THAÍS MARTINS JAJAH CARLOS

UBERLÂNDIA
2026

THAÍS MARTINS JAJAH CARLOS

DEEP LEARNING-BASED DETECTION OF CAROTID ARTERY ATHEROMAS IN PANORAMIC RADIOGRAPHS

Dissertation presented to the Graduate Program in Electrical Engineering at the School of Electrical Engineering of the Federal University of Uberlândia, as a partial requirement for the fulfillment of the Master of Science degree.

Concentration Area: Information Processing

Supervisor: Prof. Dr. Márcio José da Cunha

UBERLÂNDIA

2026

Ficha Catalográfica Online do Sistema de Bibliotecas da UFU
com dados informados pelo(a) próprio(a) autor(a).

C284 Carlos, Thaís Martins Jajah, 1989-
2026 DEEP LEARNING-BASED DETECTION OF CAROTID ARTERY ATHEROMAS IN PANORAMIC RADIOGRAPHS [recurso eletrônico] / Thaís Martins Jajah Carlos. - 2026.

Orientador: Márcio José da Cunha.
Dissertação (Mestrado) - Universidade Federal de Uberlândia, Pós-graduação em Engenharia Elétrica.
Modo de acesso: Internet.
DOI http://doi.org/10.14393/ufu.di.2026.223
Inclui bibliografia.

1. Engenharia elétrica. I. Cunha, Márcio José da,1978-, (Orient.). II. Universidade Federal de Uberlândia. Pós-graduação em Engenharia Elétrica. III. Título.

CDU: 621.3

Bibliotecários responsáveis pela estrutura de acordo com o AACR2:
Gizele Cristine Nunes do Couto - CRB6/2091
Nelson Marcos Ferreira - CRB6/3074

APPROVAL SHEET

Thaís Martins Jajah Carlos.

Deep Learning-Based Detection Of Carotid Artery Atheromas In Panoramic Radiographs.

Examining Committee Chair (Supervisor): Prof(a). Dr(a).Márcio José da Cunha

Dissertation presented to the Graduate Program in Electrical Engineering at the School of Electrical Engineering of the Federal University of Uberlândia, as a partial requirement for the fulfillment of the Master of Science degree.

Concentration Area: Electrical Engineering

Examining Committee

Full Member: Prof(a). Dr(a).Gabriela Vieira Lima
Institution: UFU

Full Member: Prof(a). Dr(a). Ana Terezinha Marques Mesquita
Institution: UFVJM

UNIVERSIDADE FEDERAL DE UBERLÂNDIA
Coordenação do Programa de Pós-Graduação em Engenharia Elétrica
Av. João Naves de Ávila, 2121, Bloco 3N - Bairro Santa Mônica, Uberlândia-MG, CEP 38400-902
Telefone: (34) 3239-4707 - www.posgrad.feelt.ufu.br - copel@ufu.br

ATA DE DEFESA - PÓS-GRADUAÇÃO

Programa de Pós-Graduação em:	Engenharia Elétrica				
Defesa de:	Dissertação de Mestrado, 819, PPGEELT				
Data:	Vinte e sete de fevereiro de dois mil e vinte e seis	Hora de início:	14:00	Hora de encerramento:	15:40
Matrícula do Discente:	12412EEL004				
Nome do Discente:	Thaís Martins Jajah Carlos				
Título do Trabalho:	Deep Learning-Based Detection of Carotid Artery Atheromas in Panoramic Radiographs				
Área de concentração:	Processamento da Informação				
Linha de pesquisa:	Metodologia e Técnicas da Computação				
Projeto de Pesquisa de vinculação:	Coordenador do projeto: Aldemir Aparecido Cavalini Junior. Título do Projeto: Sistemas de Monitoramento de Variáveis Essenciais Full Wireless (SMVE-FW). Agência financiadora: Não se aplica. Número do processo na agência financiadora: Não se aplica. Vigência do projeto: 2024-2026.				

Reuniu-se através de videoconferência, a Banca Examinadora designada pelo Colegiado do Programa de Pós-graduação em Engenharia Elétrica, assim composta:

Doutores: Gabriela Vieira Lima (UFU), Ana Terezinha Marques Mesquita (UFVJM) e Márcio José da Cunha, orientador do discente.

Iniciando os trabalhos o presidente da mesa, Dr. Márcio José da Cunha, apresentou a Comissão Examinadora e a candidata, agradeceu a presença do público, e concedeu à discente a palavra para a exposição do seu trabalho. A duração da apresentação da discente e o tempo de arguição e resposta foram conforme as normas do Programa.

A seguir o senhor presidente concedeu a palavra, pela ordem sucessivamente, aos examinadores, que passaram a arguir a candidata. Ultimada a arguição, que se desenvolveu dentro dos termos regimentais, a Banca, em sessão secreta, atribuiu o resultado final, considerando a candidata:

APROVADA.

Esta defesa faz parte dos requisitos necessários à obtenção do título de Mestre. O competente diploma será expedido após cumprimento dos demais requisitos, conforme as normas do Programa, a legislação pertinente e a regulamentação interna da UFU.

Nada mais havendo a tratar foram encerrados os trabalhos. Foi lavrada a presente ata que após lida e achada conforme, foi assinada pela Banca Examinadora.

Documento assinado eletronicamente por **Marcio José da Cunha**, **Professor(a) do Magistério Superior**, em 27/02/2026, às 15:45, conforme horário oficial de Brasília, com fundamento no art. 6º, § 1º, do Decreto nº 8.539, de 8 de outubro de 2015.

Documento assinado eletronicamente por **Gabriela Vieira Lima**, **Professor(a) do Magistério Superior**, em 27/02/2026, às 15:45, conforme horário oficial de Brasília, com fundamento no art. 6º, § 1º, do Decreto nº 8.539, de 8 de outubro de 2015.

Documento assinado eletronicamente por **Ana Terezinha Marques Mesquita**, **Usuário Externo**, em 27/02/2026, às 15:46, conforme horário oficial de Brasília, com fundamento no art. 6º, § 1º, do Decreto nº 8.539, de 8 de outubro de 2015.

A autenticidade deste documento pode ser conferida no site https://www.sei.ufu.br/sei/controlador_externo.php?acao=documento_conferir&id_orgao_acesso_externo=0, informando o código verificador **7028964** e o código CRC **9957085C**.

Referência: Processo nº 23117.005462/2026-59 SEI nº 7028964

DEDICATION

To my children, the reason behind
all my efforts, with all my love.

ACKNOWLEDGMENTS

With gratitude, I acknowledge God and all the divine beings who guide and support me throughout this journey, enlightening my path and assisting me in my continuous growth.

To Prof. Márcio, my advisor, I am grateful for the opportunities, for the support, trust, patience, and for all the knowledge shared. To Prof. Lessa, I extend my special thanks for the confidence, dedication, and essential contribution to the publication of this work.

I thank my husband for his constant presence, for encouraging me during difficult moments, and for being a partner who makes everything lighter. To my children, my strength, hope, and inexhaustible source of love: you are the reason I persist and surpass myself every day.

To my parents, for their unconditional support and for providing the foundation of who I am. To my brother, my mentor and best friend, and to my sister-in-law, for their companionship and affection. I also thank my family and friends for their encouragement, care, and for always cheering for my achievements.

Finally, I express my gratitude to the Graduate Program in Electrical Engineering at the Federal University of Uberlandia for the opportunity to receive academic training and develop scientific research. I also thank CAPES (Coordination for the Improvement of Higher Education Personnel) for the institutional support, and all professionals who are part of this university, whose daily efforts make research endeavors such as this possible.

“Every discovery begins with
the desire to serve”.
Unknown

ABSTRACT

Stroke is among the leading global causes of mortality and long-term disability resulting from neurological sequelae, with carotid atherosclerosis being one of its main etiological mechanisms. This condition often progresses silently, which limits the use of standard diagnostic examinations such as Doppler ultrasound and computed tomography angiography (CTA) for population screening. However, calcifications located at the carotid bifurcation may be incidentally visualized on panoramic dental radiographs, making this exam a potential opportunistic screening tool due to its wide availability, low cost, and coverage of the C3–C4 region, where atheromas usually manifest. In this context, the present study aimed to develop and evaluate a deep-learning-based method using the MobileNetV2 architecture for the automatic detection of carotid atheromas in panoramic radiographs, aiming to optimize the clinical value of this routine exam as a screening tool. A total of 378 publicly available and fully anonymized Regions of Interest (ROIs) were used, cropped in the carotid region (640×320 px) and divided into training, validation, and test sets (264/57/57). The images underwent normalization, grayscale channel replication, and real-time data augmentation, and class imbalance was addressed using weighted loss. The model was implemented using MobileNetV2 with pretrained weights and a two-stage training scheme, consisting of initial backbone freezing followed by partial fine-tuning (~70%), with batch normalization and dropout (0.3). On the independent test set, the model achieved robust and well-balanced performance, with 94.7% accuracy, 95.7% sensitivity, 94.1% specificity, and AUC and AUPRC values of 0.963 and 0.968, respectively, using a threshold optimized by Youden's J index. These results demonstrate strong discriminative capability and reinforce the potential of panoramic radiography combined with deep learning as an opportunistic screening tool for carotid atheromas. The proposed method shows promising performance and supports the use of routine panoramic radiographs for opportunistic screening, enabling the early referral of patients for specialized vascular evaluation and strengthening the interface between oral and systemic health.

Keywords: panoramic radiography; carotid atheroma; artificial intelligence; deep learning; MobileNetV2; stroke.

RESUMO

O acidente vascular cerebral (AVC) figura entre as principais causas globais de mortalidade e incapacidade decorrente de sequelas neurológicas, sendo a aterosclerose carotídea um dos seus principais mecanismos etiológicos. Essa condição frequentemente progride de forma silenciosa, o que limita a utilização de exames padrão, como o ultrassom Doppler e a angiotomografia computadorizada (CTA), em estratégias de triagem populacional. Entretanto, calcificações na bifurcação carotídea podem ser visualizadas incidentalmente em radiografias panorâmicas odontológicas, tornando esse exame um potencial instrumento de rastreio oportunístico, dada sua ampla disponibilidade, baixo custo e abrangência da região C3–C4, onde os ateromas tendem a se manifestar. Diante desse contexto, este trabalho teve como objetivo desenvolver e avaliar um método baseado em aprendizagem profunda, utilizando a arquitetura MobileNetV2, para a detecção automática de ateromas carotídeos em radiografias panorâmicas, de modo a otimizar o valor clínico desse exame de rotina como ferramenta de rastreio oportunístico. Foram utilizadas 378 Regiões de Interesse (ROIs) públicas e totalmente anonimizadas, recortadas na região carotídea C3–C4 (640×320 px) e divididas em conjuntos de treinamento, validação e teste (264/57/57). As imagens passaram por normalização e replicação do canal de cinza, além de aumento de dados em tempo real, e o desequilíbrio de classes foi tratado por ponderação. O modelo foi implementado com a arquitetura MobileNetV2, empregando pesos pré-treinados e um esquema de treinamento em duas etapas, com congelamento inicial do backbone seguido de fine-tuning parcial (~70%), incluindo batch normalization e dropout (0,3). No conjunto de teste independente, o modelo apresentou desempenho robusto e balanceado, alcançando acurácia de 94,7%, sensibilidade de 95,7% e especificidade de 94,1%, além de AUC de 0,963 e AUPRC de 0,968, a partir de um limiar otimizado pelo índice de Youden J. Esses resultados indicam forte capacidade discriminativa e reforçam o potencial da radiografia panorâmica, aliada a métodos de deep learning, como ferramenta de rastreio oportunístico de ateromas carotídeos. Conclui-se que o método proposto apresenta desempenho promissor e viabiliza o uso de radiografias panorâmicas de rotina como ferramenta de rastreamento oportunístico de ateromas carotídeos, favorecendo o encaminhamento preventivo de pacientes para avaliação vascular especializada e fortalecendo a interface entre saúde bucal e sistêmica.

Palavras-chave: radiografia panorâmica; ateroma carotídeo; inteligência artificial; deep learning; MobileNetV2; AVC.

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

ACC Accuracy

AI Artificial Intelligence

AUC Area Under the Receiver Operating Characteristic Curve

AUPRC Area Under the Precision–Recall Curve

BN Batch Normalization

CI Confidence Interval

CLAHE Contrast Limited Adaptive Histogram Equalization

CNN Convolutional Neural Networks

CONSORT-AI Consolidated Standards of Reporting Trials – Artificial Intelligence

CTA Computed Tomography Angiography

DL Deep Learning

FLOP Floating Point Operation

FN False Negative

FP False Positive

GAP Global Average Pooling

LVAT Latent-Space Virtual Adversarial Training

ML Machine Learning

MRA Magnetic Resonance Angiography

NPV Negative Predictive Value

PACS Picture Archiving and Communication System

PPV Positive Predictive Value

PR Precision–Recall

RIS Radiology Information System

ROC Receiver Operating Characteristic

ROI Region of Interest

SEN Sensitivity

SPE Specificity

TN True Negative

TP True Positive

TRIPOD-AI Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis – Artificial Intelligence

TTA Test Time Augmentation

XAI Explainable Artificial Intelligence

WHO World Health Organization

TABLE OF CONTENTS

1. INTRODUCTION

Stroke stands out as one of the leading global causes of mortality and long-term disability, representing a major public health challenge and imposing a considerable burden on rehabilitation systems and long-term care services (World Health Organization, 2024). In Brazil, stroke is among the principal causes of death, yielding expressive morbidity and mortality rates that directly affect quality of life and social productivity (World Health Organization, 2021). Among its etiologies, carotid atherosclerosis is one of the most relevant pathological mechanisms associated with ischemic stroke, characterized by the progressive accumulation of plaques composed of lipids, inflammatory cells, cholesterol, and calcified deposits on arterial walls (Garagoli; Masson; Barbagelata, 2024). It is estimated that 15–20% of ischemic stroke cases are directly associated with carotid atherosclerotic stenosis (Pospelov et al., 2025).

Atherosclerosis follows a silent and long-term progression, potentially evolving for years without clinically detectable manifestations (Prins et al., 2025). During this period, the risk of embolus formation increases substantially, even before significant narrowing of the arterial lumen occurs. Clinically, the investigation of carotid atherosclerosis is primarily based on advanced imaging techniques, such as Doppler ultrasound, computed tomography angiography (CTA), and magnetic resonance angiography (MRA), which are considered the gold standard for assessing stenosis and plaque morphology (David et al., 2024). However, such examinations present limitations for large-scale screening due to cost, limited availability, and frequent requirement of prior clinical suspicion, which reduces opportunities for early identification and preventive intervention in asymptomatic individuals (Pakizer et al., 2025). This anatomical relationship is illustrated in Figure 1.

Figure 1– Anatomical scheme showing the location of carotid atheroma.

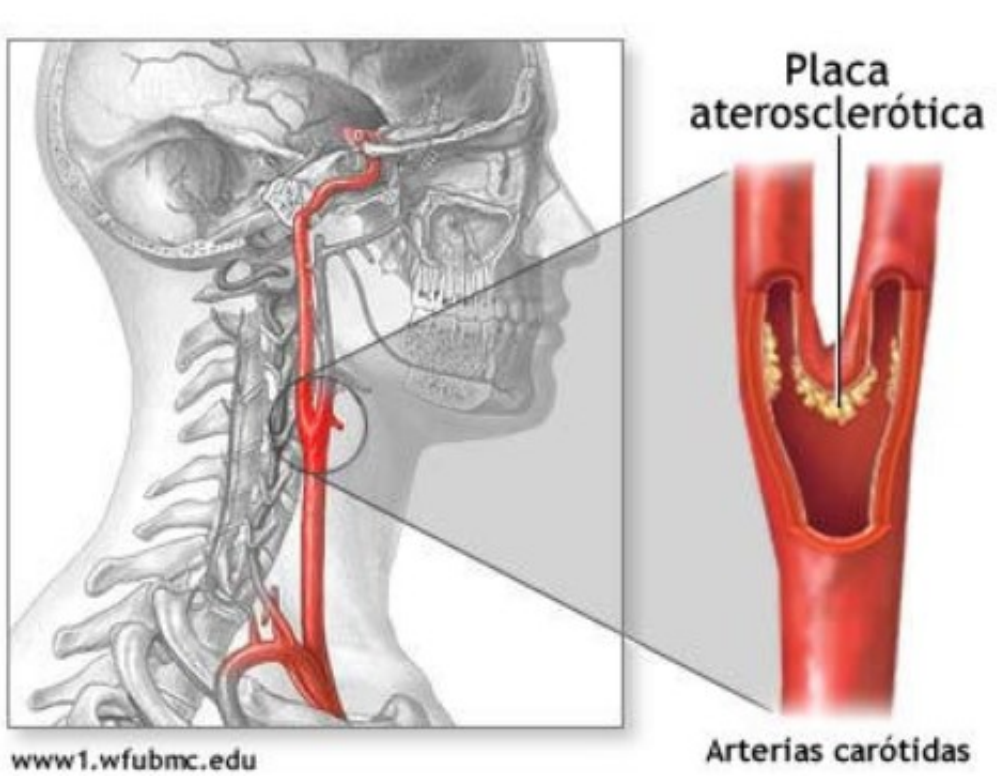

Source: Adapted from ROC Radiologia, 2023, originally attributed to Wake Forest University Health Sciences (2007).

From this perspective, the assessment of widely available imaging modalities, even when not originally designed for cardiovascular purposes, becomes highly relevant. Panoramic dental radiography—an inexpensive examination routinely performed in private and public dental services—incidentally covers the region between the cervical vertebrae C3 and C4, where the carotid bifurcation is located, a frequent site of atheroma formation (Manta et al., 2024). In this region, calcifications may appear as irregular radiopaque images and can represent potential indicators of systemic vascular risk (Schiroli et al., 2025). Thus, panoramic radiography stands out as a modality for incidental screening, enabling the identification of suggestive signs of atheroma and early referral of patients for specialized vascular evaluation (De Onofre et al., 2021).

However, manual interpretation of such findings poses significant challenges. Image assessment depends on the observer's experience and may be influenced by anatomical variability, technical artifacts, morphological similarities with other calcified structures, and differences in image acquisition, all of which increase the likelihood of diagnostic errors (Felipe et al., 2020).

Considering these limitations, Artificial Intelligence (AI), particularly deep learning (DL) techniques based on convolutional neural networks (CNNs), emerges as a promising alternative. CNNs have demonstrated superior performance in tasks involving classification and automatic pattern detection in medical images, even under conditions of low quality or high technical heterogeneity (Aggarwal et al., 2021; Zhou et al., 2021). In the context of dental imaging, recent studies have highlighted the growing applicability of deep learning methods for automated image analysis and decision support, reinforcing their potential integration into routine dental workflows (Corbella; Srinivas; Cabitza, 2021). Furthermore, optimized architectures such as MobileNetV2 offer competitive performance with reduced computational cost, favoring scalable applications that can be integrated into routine clinical systems (Alsakar et al., 2024). Consequently, the incorporation of DL into the analysis of panoramic radiographs enhances the diagnostic value of a routine dental examination, expanding its impact on secondary prevention of vascular diseases and strengthening the interface between oral and systemic health (Prados-Privado et al., 2022).

In recent years, several studies have investigated the application of artificial intelligence techniques to detect carotid artery calcifications in panoramic dental radiographs. Convolutional neural networks have demonstrated promising performance in identifying radiographic patterns associated with carotid atheromas, suggesting that automated analysis may support opportunistic screening in dental settings. Different architectures and methodological strategies have been explored, including classification and object-detection frameworks, with reported results indicating substantial potential for assisting clinicians in identifying patients at risk for vascular disease. Nevertheless, many existing studies rely on limited datasets, heterogeneous region-of-interest definitions, or computationally demanding architectures, highlighting the need for approaches that combine methodological robustness with practical feasibility for routine clinical deployment (Amitay et al., 2023; Arzani et al., 2025; Prados-Privado et al., 2022; Vinayahalingam et al., 2024).

Therefore, the general objective of this dissertation is to develop and evaluate a deep learning–based method for the automatic detection of carotid atheromas on panoramic dental radiographs, using region-of-interest images standardized around the carotid bifurcation and a convolutional neural network architecture (MobileNetV2) optimized for efficiency and robustness.

Figure 2 – Overview of the proposed deep learning pipeline for carotid atheroma detection on panoramic dental radiographs

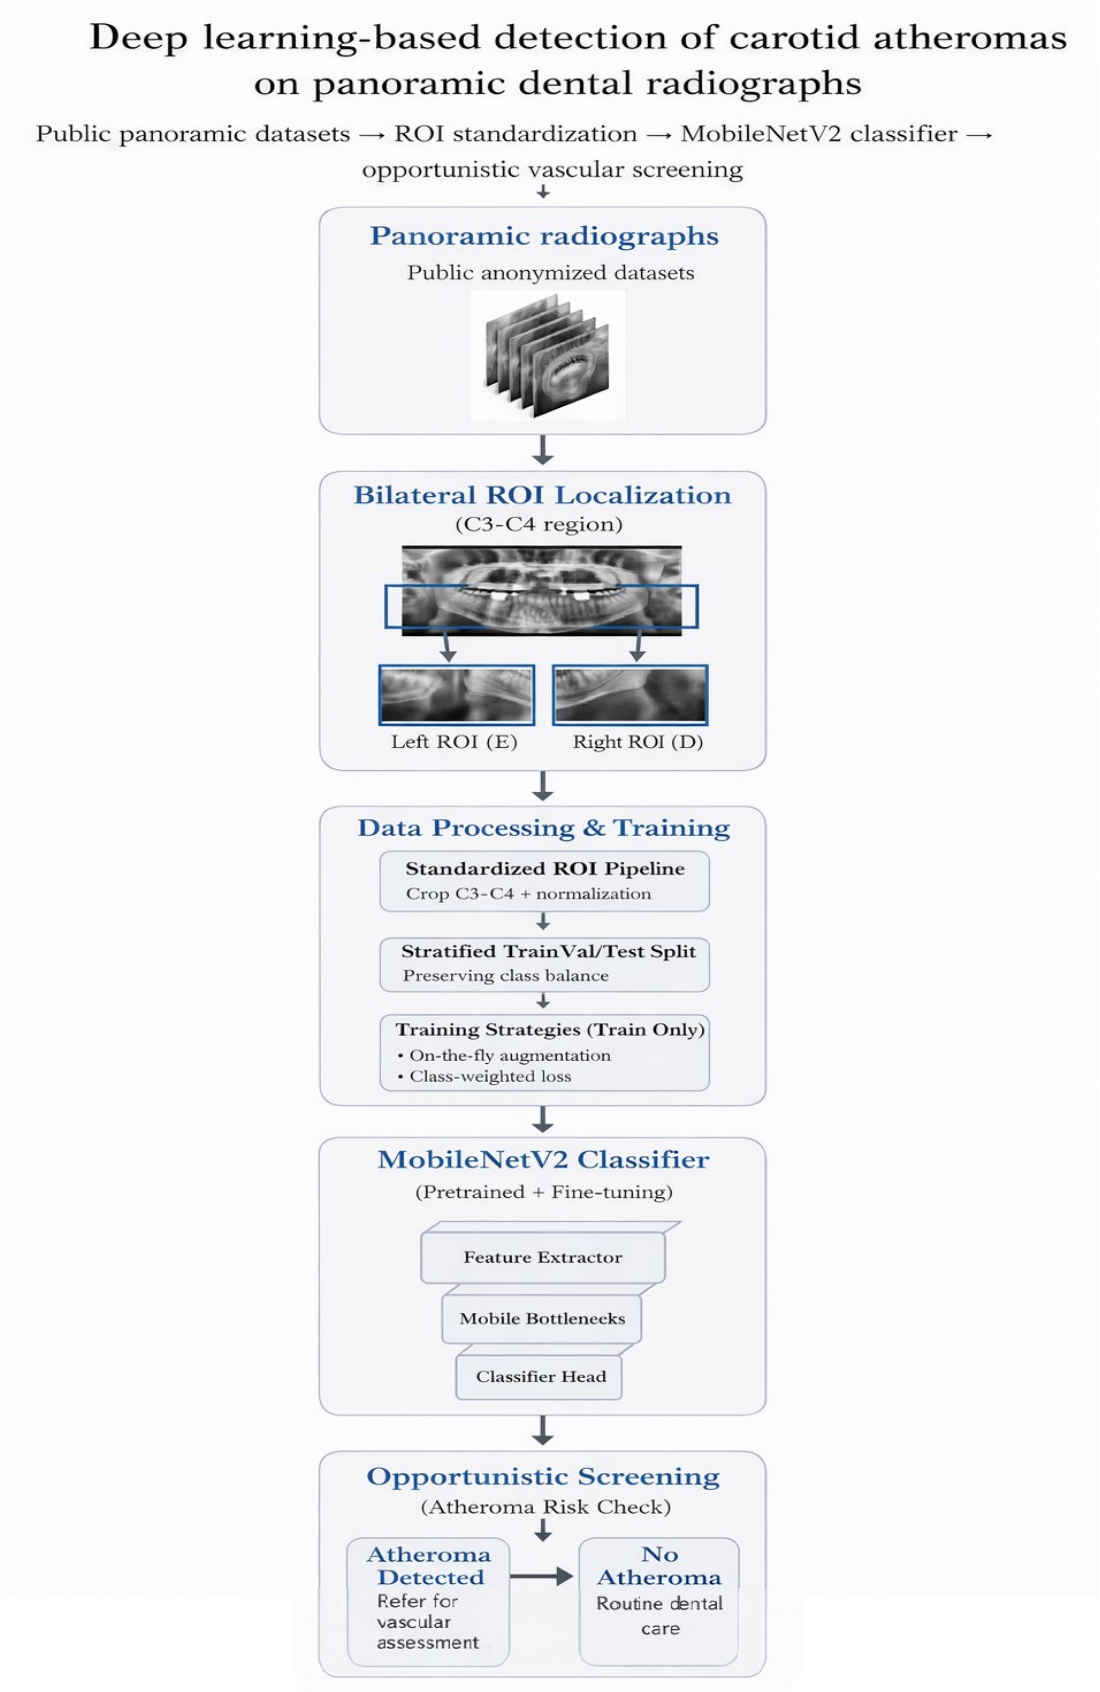

Source: Diagram elaborated by the author.

1.1 Justification

The growing prevalence of cerebrovascular diseases and the significant impact of stroke on mortality and long-term disability highlight the need for strategies that enable the early identification of modifiable risk factors (World Health Organization, 2011). Among these factors, carotid atherosclerosis stands out due to its strong association with ischemic events and its silent progression, which hinders diagnosis before acute clinical manifestations occur (Garagoli; Masson; Barbagelata, 2024). Therefore, the adoption of accessible screening methods capable of detecting early signs of atherosclerotic disease represents an important approach for reducing stroke-related morbidity and mortality (Bushnell et al., 2024).

Despite its high clinical relevance, the gold-standard examinations for carotid assessment, such as Doppler ultrasound, computed tomography angiography (CTA), and magnetic resonance angiography (MRA), present significant limitations when applied to population-based screening (David et al., 2024). These methods demand specific infrastructure, trained personnel, prior clinical indication, and incur high costs when employed preventively, especially in asymptomatic individuals (Pakizer et al., 2025). As a result, many patients at atherosclerotic risk remain unassessed, reducing opportunities for early identification and preventive intervention. In this context, it becomes essential to investigate more accessible and widely available imaging modalities capable of contributing to opportunistic screening of systemic vascular risk (Schiroli et al., 2025).

From this perspective, panoramic dental radiography emerges as a relevant alternative. Being a low-cost exam widely performed for routine dental purposes, it presents a unique opportunity to integrate dental practice with the prevention of systemic diseases (Prados-Privado et al., 2022). Its incidental coverage of the carotid bifurcation region, generally located between the C3 and C4 cervical vertebrae, enables the visualization of calcifications suggestive of atherosclerotic plaques even in asymptomatic patients (Manta et al., 2024). Population-based studies report a prevalence of carotid calcification findings on routine panoramic radiographs ranging from approximately 3% to 5% in general dental populations, increasing to around 10%–11% among high-risk individuals such as patients with diabetes mellitus, hypertension, or metabolic syndrome (Saati; Foroozandeh; Alafchi, 2020). This expands the potential

for early detection in population-level settings, especially in public health systems and routine dental consultations (De Onofre et al., 2021).

However, the effectiveness of this opportunistic screening depends on visual interpretation performed by the clinician, which introduces limitations such as interobserver variability, anatomical structures that mimic calcifications, differences in clinical experience, and heterogeneity in image quality (Felipe et al., 2020). These factors reduce diagnostic accuracy and constrain the clinical impact of screening when based solely on human observation. Thus, validating automated approaches becomes essential for ensuring greater standardization, reliability, and efficiency in identifying incidental carotid calcifications (Vinayahalingam et al., 2024).

In this scenario, Artificial Intelligence (AI) techniques, particularly those based on deep learning (DL), represent a promising solution by providing high accuracy in medical image classification tasks (Aggarwal et al., 2021), while enabling fast, scalable, and low-cost analysis (Anwar et al., 2018). Efficient architectures, such as MobileNetV2, enable the development of lightweight computational tools that can be integrated into existing clinical systems, thereby supporting their application in real-world screening settings (Alsakar et al., 2024).

In light of these considerations, the development of an automated method capable of identifying carotid atheromas in panoramic radiographs using deep learning techniques becomes pertinent. Such an approach aims to offer a scalable, accessible solution integrated into routine clinical workflows, favoring the early detection of systemic vascular risk and contributing to the reduction of stroke-related burden in the population (Arzani et al., 2025).

This dissertation is structured as follows: Chapter 2 presents the theoretical foundations related to stroke, carotid atherosclerosis, panoramic imaging, and artificial intelligence; Chapter 3 defines the general and specific objectives of the research; Chapter 4 details the materials and methods adopted; Chapter 5 presents and analyzes the experimental results; Chapter 6 discusses the findings in light of current literature; and Chapter 7 concludes the work and outlines perspectives for future studies. An overview of the proposed methodological pipeline is illustrated in Figure 2.

2. THEORETICAL FRAMEWORK

2.1 Stroke (Cerebrovascular Accident) and Carotid Atherosclerosis

2.1.1 Epidemiology and Pathophysiology

Stroke represents one of the most prevalent causes of death and long-term disability worldwide, imposing a significant public health burden that transcends economic and geographic boundaries (Abissegue et al., 2024). In Brazil, stroke mortality remains a critical indicator of social vulnerability, with evidence showing that elderly populations and those with lower socioeconomic status are disproportionately affected by the lack of timely intervention (Santos et al., 2025). The likelihood of experiencing an ischemic event increases significantly with advancing age, particularly as the prevalence of underlying vascular pathologies rises in individuals over the age of fifty (Amitay et al., 2023).

Major modifiable risk factors, including systemic arterial hypertension, smoking, diabetes mellitus, dyslipidemia, and obesity, contribute directly to the acceleration of vascular aging and atherosclerotic processes (Wei et al., 2024a). These factors often interact through complex biomechanical and inflammatory pathways, leading to endothelial dysfunction and the progressive thickening of arterial walls (Tziotziou et al., 2025). Such pathological changes promote the development of carotid plaques, which serve as the primary source of emboli in many ischemic cases, thus highlighting the critical need for diagnostic strategies that facilitate the early detection of subclinical atherosclerosis (Lee; Kim; Jeong, 2022).

In addition to their epidemiological relevance, these risk factors are directly involved in the biological mechanisms that drive atherosclerotic plaque formation and cerebrovascular events. Approximately 80% of cerebrovascular events are ischemic in origin, commonly resulting from arterial occlusion induced by thromboembolism (Abissegue et al., 2024). In this context, carotid atherosclerosis plays a central role in cerebral ischemia, accounting for an estimated 15%–20% of ischemic stroke cases (Pospelov et al., 2025).

Carotid atherosclerosis is a chronic inflammatory condition characterized by the progressive accumulation of lipoproteins, macrophages, cholesterol, and calcified crystals within the arterial intima (Garagoli; Masson; Barbagelata, 2024). It begins with endothelial dysfunction induced by oxidized low-density lipoproteins (LDL), followed

by monocyte adhesion, smooth muscle cell proliferation, extracellular matrix deposition, and the development of atheromatous plaques (Tziotziou et al., 2025). As the lesions mature, smooth muscle cells may undergo osteogenic differentiation, leading to calcium deposition and vascular calcification (Amitay et al., 2023).

These calcifications are clinically relevant because they are associated with increased plaque stiffness, reduced vascular compliance, and enhancement of local mechanical stress, which in turn contribute to plaque instability and a higher risk of embolization, especially within the carotid circulation (Lee; Kim; Jeong, 2022). Consequently, the identification of vascular calcifications serves as a potential imaging biomarker of advanced atherosclerotic burden and heightened cerebrovascular risk (Schiroli et al., 2025).

2.1.2 Carotid bifurcation anatomy and imaging relevance

The common carotid artery bifurcates into the internal and external carotid arteries typically at the level of the third to fourth cervical vertebrae (C3–C4), a region characterized by complex hemodynamic conditions including low wall shear stress, vortex formation, and flow turbulence (Saxena; Ng; Lim, 2019). These local vascular dynamics favor lipid deposition, endothelial dysfunction, and inflammatory cell recruitment , thereby predisposing the bifurcation zone to the early development and preferential accumulation of atherosclerotic plaques (Daolio et al., 2024).

Anatomically, the carotid bifurcation lies within the cervical soft tissue region encompassed by the field of view of panoramic dental radiographs. Calcified atheromatous plaques may, therefore, be incidentally visualized as irregular radiopaque nodular or linear opacities located posteroinferior to the mandibular angle and adjacent to the cervical vertebrae (Song et al., 2022). Such findings, although often overlooked in routine dental assessments (Lee; Kim; Jeong, 2022), have been shown to correlate with underlying carotid artery disease and increased risk of ischemic stroke (Tziotziou et al., 2025).

This anatomical overlap between vascular pathology and dental imaging confers panoramic radiography a unique role in opportunistic screening strategies (Zhou et al., 2021), enabling the preliminary identification of individuals who may benefit from targeted vascular assessment, particularly among asymptomatic populations undergoing routine dental care (Prados-Privado et al., 2022).

2.2 Panoramic Dental Radiography as an Opportunistic Screening Tool

2.2.1 Principles of panoramic imaging

Panoramic radiography is a two-dimensional extraoral imaging technique based on rotational tomography, in which the X-ray source and detector rotate synchronously around the patient's head, acquiring sequential projections that are reconstructed into a single panoramic image. Only anatomical structures positioned within a predetermined focal trough (image layer) are visualized with adequate sharpness, while objects located outside this layer present variable magnification, blurring, or distortion (Katı et al., 2025)(Zhu et al., 2022).

Patient positioning is a critical factor for image quality. Head tilt, rotation, and vertical displacement may alter the relationship between anatomical structures and the focal trough, causing superimposition, elongation, or compression artifacts (Mayil; Keser; Pekiner, 2014). Although panoramic imaging offers lower spatial resolution compared with dedicated maxillofacial and vascular imaging modalities, its extended field of view enables visualization of adjacent cervical soft tissues, including the anatomical region of the carotid artery bifurcation (Katı et al., 2025).

2.2.2 Visualization of the Carotid Region and Radiographic Appearance of Calcifications

Panoramic dental radiography incidentally encompasses the anatomical region corresponding to the third and fourth cervical vertebrae (C3–C4), where the common carotid artery bifurcates into the internal and external carotid arteries. This region lies posterior and inferior to the mandibular angle within the cervical soft tissue field of view. Because atherosclerotic changes preferentially develop at the carotid bifurcation due to local hemodynamic factors, this portion of the vascular anatomy represents a clinically relevant site for opportunistic imaging assessment (Kats et al., 2019).

Although panoramic radiography was not designed for cardiovascular evaluation, its consistent coverage of this anatomical zone enables the incidental visualization of calcified atherosclerotic plaques, even in patients undergoing radiographic examination solely for dental diagnostic purposes (Kats et al., 2019; Nelson; Vaddi; Tadinada, 2024) .

When present, these calcifications typically appear on panoramic radiographs as heterogeneous radiopaque nodular, linear, or irregular foci located posteroinferior to the

mandibular angle and adjacent to the cervical vertebrae (Saati; Foroozandeh; Alafchi, 2020). Their morphology is frequently variable, reflecting differing stages of plaque calcification and mineral deposition. These radiographic findings do not allow for quantification of luminal stenosis or assessment of plaque composition but serve as surrogate markers of underlying vascular pathology associated with increased cerebrovascular risk (Janiszewska-Olszowska et al., 2022). An example of this radiographic presentation is illustrated in Figure 3 (Er, 2023).

Figure 3. Example of panoramic dental radiograph showing a calcified carotid atheroma (blue arrow).

Source: DENTEX Challenge 2023 dataset (Er, 2023).

Population-based studies report that the prevalence of carotid calcification findings on routine panoramic radiographs ranges from approximately 3% to 5% in general dental populations (Saati; Foroozandeh; Alafchi, 2020), increasing to upwards of 10%–11% among high-risk groups such as individuals with diabetes mellitus, hypertension, or metabolic syndrome. Although these radiographic findings are not diagnostic equivalents of vascular imaging studies, they have been consistently recognized as valuable alerts for referral to targeted Doppler ultrasound examination for definitive vascular assessment (Kats et al., 2019; Saati; Foroozandeh; Alafchi, 2020).

2.2.3 Diagnostic Limitations and Interpretation Challenges

The gold-standard diagnostic methods for evaluating carotid artery disease include Doppler ultrasound, computed tomography angiography (CTA), and magnetic resonance angiography (MRA). Doppler ultrasound is the most widely used noninvasive

technique for initial vascular assessment, providing real-time measurements of blood-flow velocity and enabling indirect estimation of luminal stenosis. CTA and MRA offer high-resolution anatomical visualization of the carotid arteries, allowing precise assessment of plaque morphology, degree of narrowing, vascular wall composition, and potential complications. These modalities are commonly employed in preoperative evaluation or in cases requiring advanced diagnostic clarification (Daolio et al., 2024).

Despite their high diagnostic accuracy, these examinations present important limitations for population-based or opportunistic screening. They require dedicated equipment, trained operators, relatively high operational costs, and are generally indicated only after clinical suspicion has been established. Consequently, their application is largely restricted to symptomatic individuals or patients with known cardiovascular risk, limiting early detection in asymptomatic populations. In contrast, panoramic radiography is routinely acquired for dental purposes, offering a unique opportunity for incidental identification of vascular calcifications without additional radiation exposure, financial cost, or patient discomfort (Dashti et al., 2025).

However, the reliable identification of carotid calcifications on panoramic radiographs remains challenging and is subject to significant interobserver variability. Several anatomical and pathological structures may mimic carotid atherosclerotic calcifications, including tonsilloliths, the stylohyoid ligament complex, calcified triticeal cartilage, sialoliths, phleboliths, and calcified lymph nodes (Saati; Foroozandeh; Alafchi, 2020). Additionally, image artifacts, overlapping anatomic shadows, patient positioning errors, and variations in acquisition parameters further complicate interpretation (Mayil; Keser; Pekiner, 2014).

These confounding factors can result in both false-positive and false-negative assessments when image evaluation relies solely on manual inspection by dentists or oral radiologists, thereby limiting the effectiveness of opportunistic screening based on human interpretation alone. This scenario highlights the need for standardized, reproducible, and scalable analysis methods, providing strong justification for the application of automated artificial intelligence–based systems to assist clinicians in detecting subtle carotid calcifications in panoramic imaging (Ikeda et al., 2024; Kats et al., 2019).

2.3 Artificial Intelligence (AI) and Deep Learning in Medical Image Analysis

2.3.1 AI, ML, and DL: definitions

Artificial Intelligence (AI) is broadly defined as the field of computer science dedicated to the development of computational systems capable of performing tasks that typically require human intelligence, including learning, reasoning, pattern recognition, decision-making, and visual perception. Within this domain, Machine Learning (ML) refers to algorithms designed to learn patterns directly from data without explicit programming, progressively optimizing their performance based on experience (Johnson et al., 2021).

Deep Learning (DL) constitutes a specialized branch of ML founded on artificial neural networks composed of multiple interconnected layers (LeCun; Bengio; Hinton, 2015). These architectures enable hierarchical representation learning, allowing computational models to automatically extract increasingly complex and abstract features from high-dimensional input data. DL has become particularly prominent in medical image analysis, where large volumes of heterogeneous visual data require robust feature extraction frameworks capable of handling noise, variation in acquisition parameters, and anatomical complexity (Wei et al., 2024b). A simplified representation of a deep neural network architecture is illustrated in Figure 4.

Figure 4 - Schematic representation of a deep neural network composed of an input layer, multiple hidden layers, and an output layer, where each node represents an artificial neuron.

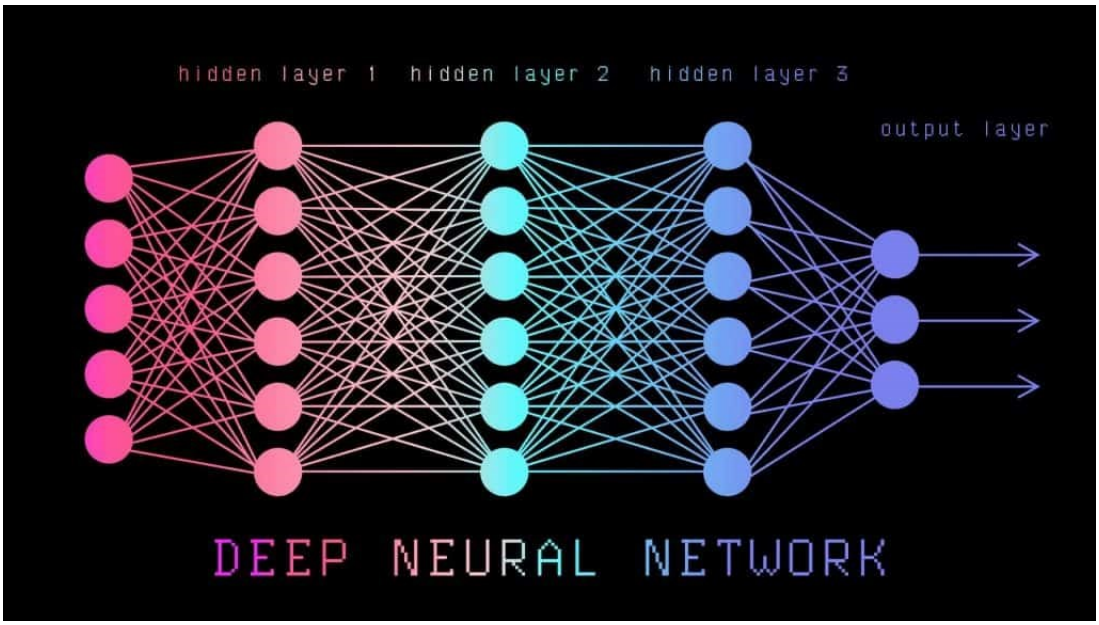

Source: Cybermagician, Shutterstock, 2023.

2.3.2 Fundamentals of Convolutional Neural Networks (CNNs)

Convolutional Neural Networks (CNNs) represent the most widely adopted DL architecture for two-dimensional image processing. Their core structure is composed of convolutional layers that apply trainable kernels (filters) across localized regions of input images to detect spatial features such as edges, shapes, textures, and contrast gradients. Following convolution operations, nonlinear activation functions introduce representational flexibility, allowing the network to model complex image patterns (LeCun; Bengio; Hinton, 2015).

Pooling layers progressively reduce spatial resolution and computational complexity while preserving salient feature information and improving translational invariance. Finally, fully connected layers integrate the extracted features and map them into output classes or quantitative predictions. This hierarchical architecture permits the automated extraction of discriminative features directly from raw image data, eliminating dependence on handcrafted descriptors and enabling robust performance across varying imaging conditions (Goodfellow; Bengio; Courville, 2016).

2.3.3 DL applications in medical imaging

The use of Artificial Intelligence in medicine has expanded exponentially over the past decade, particularly in applications centered on automated interpretation of medical images (Esteva et al., 2021; Johnson et al., 2021). CNN-based models have demonstrated high diagnostic performance across a broad range of imaging modalities, including radiography, computed tomography, magnetic resonance imaging, and ultrasound. These systems have been successfully deployed in tasks such as tumor detection, lesion segmentation, organ delineation, and disease classification.

In the vascular domain, DL approaches have shown effectiveness in automated assessment of carotid stenosis using CTA (Omarov et al., 2025), real-time plaque recognition on carotid ultrasound images (Wei et al., 2024b), and risk stratification based on plaque morphology. Similarly, in the dental field, CNNs have been applied to the detection of maxillofacial anatomical landmarks, identification of odontogenic lesions, diagnosis of dental caries, and recognition of calcified structures on panoramic radiographs (Kats et al., 2019; Nelson; Vaddi; Tadinada, 2024).

In addition to achieving high accuracy, CNN-based solutions may be optimized using strategies such as transfer learning, model fine-tuning, intensity normalization,

and data augmentation, enabling improved generalization even when training datasets are limited (Omarov et al., 2025). Furthermore, architectures with reduced computational complexity support efficient deployment in conventional clinical software environments, broadening access to automated image analysis tools beyond large tertiary medical centers. This scalability is particularly relevant for opportunistic screening scenarios in dental practice, where specialist availability may be limited and real-time decision support can facilitate early risk identification and appropriate referral pathways (Janiszewska-Olszowska et al., 2022; Saati; Foroozandeh; Alafchi, 2020).

2.4 CNNs in Dentistry and Selection of the MobileNetV2 Architecture

2.4.1 Deep Learning Applications in Dentistry and Panoramic Radiography

Several studies have validated the applicability of deep learning approaches, particularly convolutional neural networks (CNNs), in dental imaging. CNN-based systems have demonstrated strong performance in tasks such as detection of bone lesions, dental caries classification, mandibular fracture identification, segmentation of complex anatomical structures, and recognition of soft-tissue calcifications on panoramic radiographs and cone-beam computed tomography images (Saati; Foroozandeh; Alafchi, 2020; Vinayahalingam et al., 2024). These applications highlight the potential of automated image analysis to support dental diagnostics, reduce interpretation variability, and enhance clinical decision-making (Kats et al., 2019).

The extension of DL methodologies to panoramic dental radiography is especially promising due to the widespread availability of this imaging modality and its potential role in opportunistic screening initiatives. Recent advances have demonstrated that CNNs can effectively detect subtle radiographic patterns even in heterogeneous image datasets, thereby expanding their applicability beyond traditional dental pathology to systemic disease indicators visible on dental examinations (Janiszewska-Olszowska et al., 2022).

Beyond dental imaging, deep learning techniques have also been widely applied to the assessment of carotid atherosclerosis using dedicated vascular imaging modalities. CNN-based models have demonstrated high accuracy in identifying and segmenting carotid plaques on computed tomography angiography (CTA), enabling automated quantification of luminal stenosis and evaluation of plaque morphology (Omarov et al., 2025). Similarly, ultrasound-based DL systems have been implemented

for real-time plaque detection and classification in carotid Doppler examinations, achieving robust performance under controlled imaging conditions. While these results confirm the clinical utility of DL for vascular diagnostics, their dependence on specialized equipment and trained operators limits their scalability for population-based screening, particularly in asymptomatic individuals (Wei et al., 2024b).

In contrast, only a limited number of studies have investigated the detection of carotid calcifications specifically on panoramic dental radiographs. Nelson; Vaddi; Tadinada, 2024 explored the use of CNNs to identify external carotid artery calcifications, while Vinayahalingam et al., 2024 applied earlier architectures to detect cervical calcifications. Collectively, these studies demonstrated the feasibility of automated detection in panoramic radiography and provided proof-of-concept evidence supporting the value of AI-assisted opportunistic screening.

Nevertheless, the current literature remains constrained by several methodological limitations. Most studies rely on relatively small, monocentric datasets, exhibit heterogeneity in region-of-interest (ROI) selection strategies, and lack standardized preprocessing pipelines. In many reports, curated subsets of images are manually selected for analysis rather than representative full clinical datasets, potentially introducing selection bias and inflating performance metrics (Janiszewska-Olszowska et al., 2022). Furthermore, inconsistent evaluation protocols and limited independent validation restrict the reproducibility and generalizability of the reported findings (Kelly et al., 2019). These gaps highlight the need for more robust, standardized DL approaches explicitly designed for real-world panoramic imaging conditions (Janiszewska-Olszowska et al., 2022).

2.4.2 Justification for using MobileNetV2

Traditional CNN architectures such as VGG and ResNet (He et al., 2016; Simonyan; Zisserman, 2015), although highly effective in high-resource environments, typically involve millions of parameters and substantial computational demands, requiring advanced hardware infrastructure and large training datasets. These constraints pose challenges for clinical integration, particularly in settings focused on opportunistic screening where rapid inference, cost-efficiency, and scalability are essential requirements (Esteva et al., 2021; Omarov et al., 2025).

MobileNetV2 was designed to address these limitations. The architecture employs depthwise separable convolutions combined with inverted residual blocks and linear bottleneck layers, substantially reducing parameter count and computational complexity while preserving strong representational capacity. Consequently, MobileNetV2 offers an efficient balance between accuracy and computational cost, making it particularly suitable for applications involving small and heterogeneous datasets—such as panoramic radiographic collections—and enabling deployment in conventional clinical computing environments without reliance on specialized hardware (Sandler et al., 2019).

Furthermore, the compatibility of MobileNetV2 with transfer learning frameworks facilitates rapid adaptation to domain-specific tasks, enhancing convergence stability and performance reliability during training (Sandler et al., 2019). These attributes support efficient model development even when large labeled datasets are not available (Esteva et al., 2021).

In this context, transfer learning refers to the process of leveraging knowledge obtained from large-scale pretrained models and adapting it to a new, domain-specific task (Alzubaidi et al., 2021). Pretrained weights derived from datasets such as ImageNet (Russakovsky et al., 2015) capture generalized low-level visual features including edges, textures, and contrast patterns, which are transferable to medical imaging domains (Yosinski et al., 2014). By fine-tuning these pretrained representations for the detection of carotid atheromas on panoramic radiographs, the network benefits from accelerated training convergence, improved generalization capability, and reduced risk of overfitting, particularly in scenarios characterized by limited annotated datasets (Alzubaidi et al., 2021; Yosinski et al., 2014).

Therefore, the selection of MobileNetV2 is justified not only by its technical merits but also by its alignment with real-world clinical requirements, including low computational cost, fast inference times, ease of integration into existing radiographic workflows, and scalability for community-based screening initiatives (Esteva et al., 2021; Sandler et al., 2019). Within the context of panoramic radiography, this architecture represents a viable and promising solution for automated detection of carotid atheromas, strengthening the potential contribution of dental imaging to secondary prevention strategies for systemic vascular disease (Janiszewska-Olszowska et al., 2022).

3. OBJECTIVES

3.1 General Objectives

To develop and validate a deep learning–based method for the automatic detection of calcifications consistent with carotid atheromas in standardized Regions of Interest (ROIs) extracted from panoramic dental radiographs, using the MobileNetV2 architecture, in order to support opportunistic screening of vascular risk in routine dental practice and facilitate the early referral of asymptomatic patients for further clinical evaluation.

3.2 Specific Objectives

- To perform curation and standardization of a public dataset of panoramic radiographs, extracting Regions of Interest (ROIs) centered on the carotid bifurcation (C3–C4) and ensuring anonymization and anatomical consistency.
- To implement a robust preprocessing pipeline, including intensity normalization, channel replication, standardized ROI extraction (640×320 px), and real-time data augmentation (on-the-fly), in order to mitigate overfitting resulting from small and heterogeneous datasets.
- To apply a two-stage training strategy, consisting of an initial backbone freeze (training only the classifier head), followed by partial fine-tuning of the upper blocks of MobileNetV2, using transfer learning and class imbalance correction strategies.
- To evaluate the model's performance using an independent test set, employing clinically relevant metrics such as Accuracy, Sensitivity (SEN), Specificity (SPE), AUC, and AUPRC, under two operational points: (i) optimal thresholding according to the Youden J index (maximizing class balance), and (ii) a sensitivity-oriented thresholding strategy (≈ 0.80) for screening purposes.

To compare the performance and computational efficiency of the proposed method with previous studies reported in the literature, assessing its feasibility for clinical integration in opportunistic screening workflows.

4 MATERIALS AND METHODS

4.1 Type of Study

This study is classified as a methodological, experimental, and quantitative investigation with a technological development focus, aimed at proposing, implementing, and validating a deep learning–based computational model for the automated detection of carotid atheromas in panoramic dental radiographs.

The research falls within the scope of applied research, as it seeks to address a specific clinical problem—opportunistic screening of carotid calcifications—through the development and evaluation of a computational system designed for potential application in real-world healthcare settings. No direct intervention involving human subjects was performed. All analyses were conducted using publicly available, previously acquired, and fully anonymized image datasets, characterizing this work as a secondary data analysis study.

The experimental approach comprised data curation, standardized extraction of regions of interest (ROIs), image preprocessing, and implementation of a supervised training pipeline using transfer learning and fine-tuning of the selected CNN architecture. Model performance evaluation was carried out on an independent test set using classification metrics commonly applied to computer-aided diagnosis studies (Esteva et al., 2021), including accuracy, sensitivity, specificity, area under the ROC curve (AUC), and area under the precision–recall curve (AUPRC) (Saito; Rehmsmeier, 2015). Additionally, stability analyses were performed using bootstrap resampling and decision-threshold optimization strategies (Efron, B. and Tibshirani, R.J., 1994).

This study presents a non-probabilistic design, without population-representative sampling, and is therefore characterized as a computational proof-of-concept experiment focused on assessing the technical feasibility, discriminative performance, and potential clinical utility of convolutional neural networks for the detection of carotid atheromas in panoramic radiographic images.

From a temporal perspective, the study is classified as cross-sectional, as the analysis involved images obtained at different time points without longitudinal follow-up of patients. All experimental procedures were conducted exclusively within a controlled computational environment, with no prospective clinical data collection.

4.2 Ethical Aspects

The present study was conducted in compliance with ethical principles governing research involving human subjects, as established by the Declaration of Helsinki (World Medical Association, 2025) and by the guidelines of the Brazilian National Health Council (Resolutions CNS No. 466/2012 and No. 510/2016) (Conselho Nacional de Saúde, 2012, 2016). The research did not involve prospective data collection or direct interaction with human participants, being based exclusively on the secondary analysis of previously acquired, publicly available, and fully anonymized image datasets.

Both datasets employed in this study were obtained from open-access repositories and had been previously approved by their respective institutional ethics committees at the time of original data collection. In the case of the dataset proposed by (Mureșanu; Hedeșiu; Iacob, 2025), ethical approval was granted under reference number 117/04.06.2024, with a waiver of informed consent due to its retrospective design. The DENTEX Challenge dataset (Er, 2023) was released for research purposes following anonymization and confidentiality procedures defined by the contributing institutions.

No personally identifiable information—such as names, medical record numbers, dates of birth, or other sensitive identifiers—was accessible at any stage of this study. Therefore, this investigation is classified as research involving the secondary use of fully anonymized public-domain data and meets the criteria for exemption from submission to the Brazilian Research Ethics Committee system (CEP/CONEP), as established under Article 1, item III, of Resolution CNS No. 510/2016 (Conselho Nacional de Saúde, 2016).

All data processing and analysis procedures were conducted in a secure computational environment, in strict adherence to principles of confidentiality, scientific integrity, transparency, and reproducibility. The ethical use of open-access data for AI training followed established frameworks for data privacy and algorithmic transparency (Geis et al., 2019).

4.3 Dataset Description

4.3.1 Source of the Public Datasets

The radiographic images analyzed in this study were obtained from two international publicly available datasets released for scientific research under a Creative Commons Attribution 4.0 (CC BY 4.0) license. All images correspond to digitally acquired panoramic radiographs obtained directly from digital imaging systems, rather than scans of conventional film radiographs.

The first dataset, proposed by Mureșanu, Hedeșiu, and Iacob, consists of panoramic radiographs retrospectively collected at the Department of Oral and Maxillofacial Surgery and Radiology of the Iuliu Hațieganu University of Medicine and Pharmacy, Cluj-Napoca, Romania. Image acquisition was performed between April 2021 and December 2023 using a Vatech PCH-2500 panoramic imaging system (Vatech, Hwaseong, South Korea). The dataset includes adult patients with permanent dentition, and all images were stored in JPEG format following de-identification prior to public release (Mureșanu; Hedeșiu; Iacob, 2025).

The second dataset corresponds to the DENTEX Challenge 2023, organized within the scope of the MICCAI international conference. This multicentric dataset comprises panoramic radiographs acquired at three different institutions using heterogeneous imaging devices and routine clinical acquisition protocols, reflecting real-world variability in dental imaging. The dataset includes examinations from patients aged 12 years or older, randomly sampled from hospital databases, with all images anonymized prior to their release for academic and scientific purposes (Er, 2023).

4.3.2 Imaging Characteristics

All images consisted of digital panoramic dental radiographs acquired under routine clinical conditions. The first dataset employed a standardized device (Vatech PCH-2500), whereas the DENTEX dataset incorporated radiographs obtained from multiple panoramic systems with varied manufacturers and acquisition parameters, including differences in exposure settings, spatial resolution, field of view, and post-processing techniques.

This heterogeneity resulted in variations in image quality, including differences in contrast, sharpness, anatomical coverage, and the presence of overlapping structures.

Such variability was considered desirable for model development, as it better reflects real-world clinical scenarios and contributes to improving the generalization capability of the proposed deep learning framework.

All images were stored in two-dimensional digital format without identifiable metadata and were subsequently processed for standardized region-of-interest extraction and model training.

4.3.3 Inclusion and Exclusion Criteria

For inclusion in the study, radiographs were required to allow adequate visualization of the carotid artery bifurcation, typically projected around the cervical vertebral levels C3–C4.

The inclusion criteria were:

- Digital panoramic radiographs of sufficient diagnostic quality;
- Clear visualization of the cervical paravertebral region corresponding to the carotid anatomy;
- Absence of major obstructions preventing reliable ROI extraction;
- Clearly identifiable presence or absence of calcifications compatible with carotid atheromas based on visual assessment.

The exclusion criteria were(Saati; Foroozandeh; Alafchi, 2020):

- Images containing prominent graphic artifacts such as arrows, text overlays, tracings, or clinical markings;
- Radiographs affected by severe motion blur, acquisition errors, or inadequate sharpness compromising interpretation;
- Examinations featuring extensive metallic restorations, tumors, or large lesions obscuring the C3–C4 anatomical region;
- Duplicate images between datasets;
- Radiographs failing to adequately capture the anatomical region of interest.

4.3.4 Dataset Composition and Class Distribution

A total of 3,960 panoramic radiographs were initially evaluated across the two datasets. After applying the inclusion and exclusion criteria and dataset curation to eliminate duplicates and unsuitable images, 153 radiographs presenting calcifications compatible with carotid atheromas were identified. From the remaining negative

examinations, 225 radiographs without atheroma-related calcifications were randomly selected in order to avoid class imbalance during model training. Consequently, the final dataset used in this study comprised 378 panoramic radiographs, from which standardized regions of interest (ROIs) were extracted for model training.

Each ROI corresponded to a standardized area of 640 × 320 pixels, manually extracted from the lateral cervical region near vertebral levels C3 and C4, where visible carotid atheromas most frequently appear on panoramic radiographs.

The complete ROI set was stratified and fixedly divided into three subsets:

- **Training set:** 264 ROIs (≈70%)
- **Validation set:** 57 ROIs (≈15%)
- **Test set:** 57 ROIs (≈15%)

This partitioning ensured maintenance of similar class proportions across all subsets, enabling consistent evaluation of model generalization across experimental phases.

Following dataset partitioning, the ROIs were labeled into two diagnostic categories:

- **Class 0 (negative):** absence of radiopaque calcifications compatible with carotid atheromas;
- **Class 1 (positive):** presence of radiopaque areas suggestive of carotid atheromas at the carotid bifurcation.

The final class distribution was as follows:

Entire dataset (378 ROIs)

- 225 negative (59.5%)
- 153 positive (40.5%)

Training set (264 ROIs)

- 157 negative
- 107 positive

Validation set (57 ROIs)

- 34 negative
- 23 positive

Test set (57 ROIs)

- 34 negative
- 23 positive

A slight class imbalance toward the negative category was observed. To mitigate its potential effects on the learning process, class weighting techniques were applied during training alongside on-the-fly data augmentation, which was restricted exclusively to the training subset, as detailed in the subsequent methodological sections.

4.4 Extraction of Regions of Interest (ROIs)

4.4.1 Anatomical Identification of the C3–C4 Region

The extraction of Regions of Interest (ROIs) was targeted to the anatomical area corresponding to the carotid artery bifurcation, which is classically located at the level of the C3 and C4 vertebral bodies (Saati; Foroozandeh; Alafchi, 2020). This region is widely recognized in the literature as the most common site for the radiographic visualization of atherosclerotic calcifications on panoramic dental examinations (Janiszewska-Olszowska et al., 2022). The anatomical localization of the C3–C4 region and the extraction of bilateral ROIs are illustrated in Figure 5.

Figure 5 - Bilateral ROI localization in the C3–C4 region on panoramic radiography.

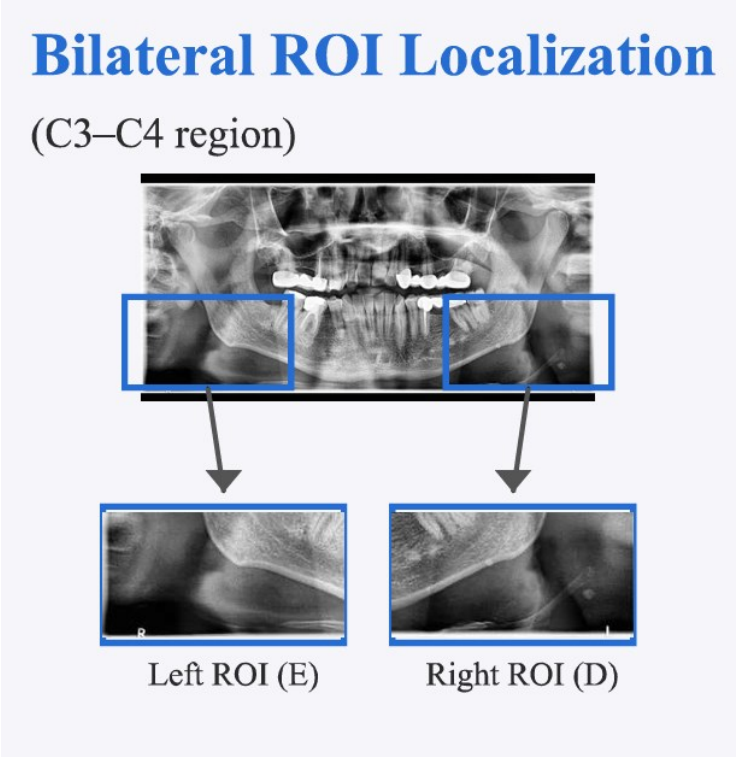

Source: Author, 2026.

On panoramic radiographs, the C3–C4 region is positioned posterior to the mandibular ramus and inferior to the mandibular angle, projecting vertically near the junction between the third and fourth cervical vertebrae. Identification of this anatomical landmark was performed through visual inspection based on adjacent osseous reference structures (Mayil; Keser; Pekiner, 2014), including:

- Contours of the mandibular ramus and angle;
- Cervical vertebral shadows;
- Intervertebral spaces corresponding to the C3–C4 levels;
- Soft-tissue regions lateral to the cervical vertebrae.

These anatomical references enabled consistent localization of the target field of interest despite the heterogeneity in image quality, contrast, and patient positioning observed across the different public datasets used in this study.

4.4.2 Standardization of the ROI Extraction Procedure

Following anatomical identification of the C3–C4 region, a standardized manual cropping procedure was applied to each panoramic radiograph, aiming to isolate exclusively the lateral cervical area corresponding to the projection of the carotid artery bifurcation.

The extraction protocol was systematically conducted according to the following criteria:

- Selection of the lateral area adjacent to the cervical spine corresponding to the carotid bifurcation region;
- Centering of the crop around the approximate C3–C4 plane, ensuring inclusion of the full area of potential atheromatous calcification;
- Exclusion of structures irrelevant to the analysis (upper and lower dental arches, maxillae, maxillary sinuses, and most mandibular structures) to reduce visual noise;
- Preservation of the spatial proportions between crops to avoid geometric distortions that might negatively affect model learning.

When both cervical sides were visible, the side presenting greater clarity of the carotid region was selected for the final ROI. In cases of anatomical uncertainty or insufficient image quality, the radiograph was excluded to preserve dataset reliability.

The standardization of ROI extraction aimed to:

- Minimize spatial variability among samples;
- Increase the signal-to-noise ratio of the input data;
- Ensure exclusive anatomical focus on the target region;
- Improve the generalization capability of the deep learning model.

4.4.3 Final ROI Dimensions (640 × 320 px)

All extracted ROIs were resized to fixed final dimensions of **640 × 320 pixels** (width × height), preserving the horizontal aspect ratio characteristic of the cervical anatomy observed on panoramic radiographs.

The choice of this resolution represented a compromise between:

- Preservation of clinically relevant radiographic details, including irregular borders, density patterns, and microcalcifications compatible with atherosclerotic plaques;
- Computational feasibility, enabling stable model training without excessive GPU memory consumption;
- Input standardization to ensure compatibility with the MobileNetV2 architecture adopted in this study.

Although the original panoramic images were provided in grayscale format (single channel), the extracted ROIs were subsequently converted to a three-channel format by replicating the grayscale channel, in order to meet the input requirements of the MobileNetV2 model pretrained on the ImageNet dataset—a procedure detailed further in the preprocessing section.

The final ROI format ensured the retention of sufficient structural information for discriminating between positive and negative carotid atheroma cases, supporting consistent data preprocessing and reliable model training, validation, and testing.

4.5 Image Preprocessing

Image preprocessing was performed to standardize input data, reduce variability arising from differences in radiographic acquisition protocols, and enhance the robustness of the deep learning model against the intrinsic heterogeneity of panoramic images. All processing steps were applied exclusively to the Regions of Interest (ROIs) extracted as described in Section 4.4.

4.5.1 Intensity Normalization

Initially, all ROIs were maintained in grayscale format and subjected to pixel intensity normalization, with original values rescaled to the continuous range [0, 1].

This step aimed to:

- Reduce global variability resulting from differences in radiographic exposure, contrast, and equipment post-processing;
- Standardize intensity distributions across all samples;
- Improve optimization stability during network training by preventing gradient saturation effects.

Additional image enhancement techniques, including contrast equalization or smoothing filters, were deliberately avoided in order to preserve native image textures and microstructural details considered relevant for detecting subtle calcifications associated with carotid atheromas.

Images presenting excessive noise or severe interpretation impairment had already been excluded during the dataset curation stage, as described in Section 4.3.

4.5.2 Channel Replication

The MobileNetV2 architecture adopted in this study was initialized with pretrained weights from the ImageNet dataset, which requires three-channel (RGB) image input.

Because panoramic radiographs are inherently grayscale (single-channel), channel replication was employed, whereby the original intensity map was duplicated to populate all three channels of the final input image.

This strategy enabled:

- Direct compatibility with the pretrained architecture while preserving the benefits of transfer learning;
- Avoidance of artificial modification of visual information, as no chromatic transformations or synthetic color mappings were introduced;
- Preservation of the integrity of the original radiographic content, ensuring that feature learning was driven exclusively by genuine structural and intensity variations present in the ROIs.

Channel replication is a well-established approach in medical deep learning studies using architectures originally designed for color imagery and does not introduce relevant bias into feature extraction.

4.5.3 On-the-Fly Data Augmentation

To increase the diversity of the training dataset and reduce overfitting associated with the relatively limited sample size, real-time data augmentation strategies were applied exclusively during model training. A conceptual illustration of the applied augmentation strategies is presented in Figure 6.

Figure 6 - Conceptual illustration of data augmentation strategies applied during model training.

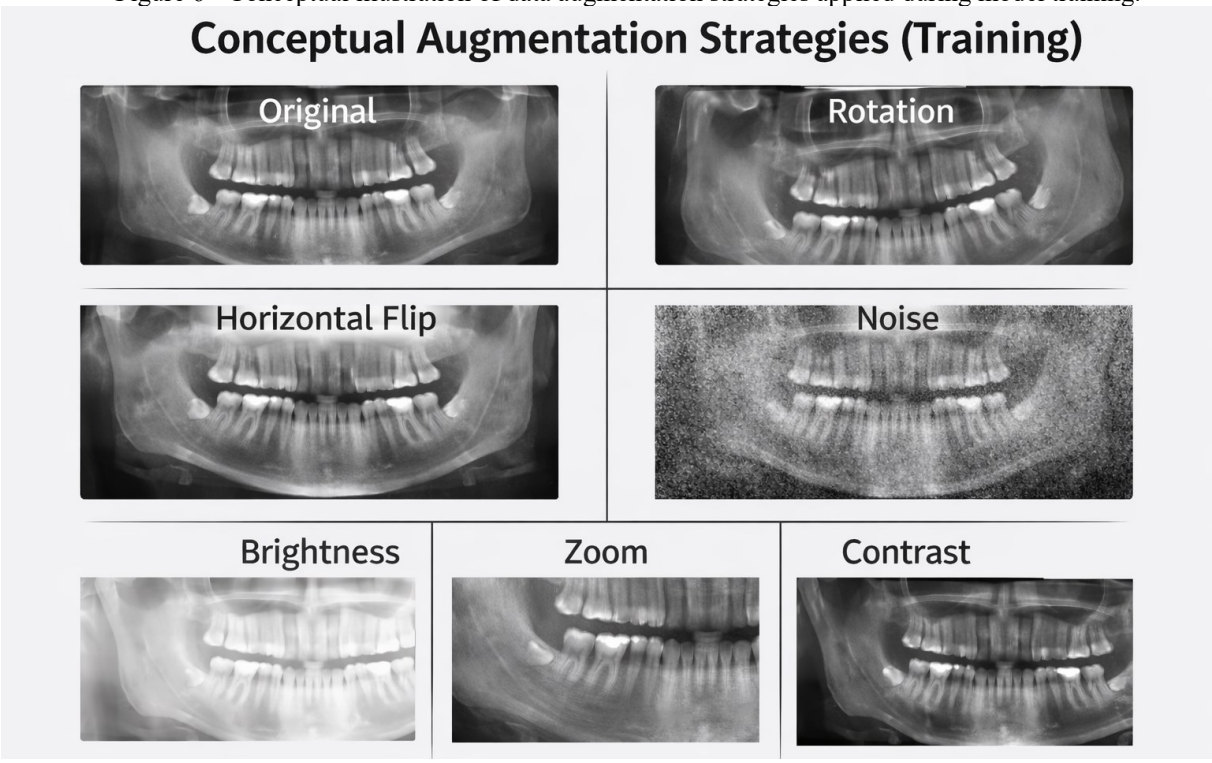

Source: Author (AI-generated illustration using ChatGPT), 2026.

Augmentation was implemented using Keras framework layers and performed dynamically on each mini-batch, without permanent storage of transformed images. The applied transformations included:

- Random horizontal flipping;
- Mild rotations of up to approximately ±5°;
- Horizontal and vertical translations of up to 5% of image dimensions;
- Zoom variations of up to 10%.

These operations simulate realistic variations in patient positioning, acquisition angles, and anatomical framing commonly observed in routine clinical practice, while

preserving the morphological integrity of the C3–C4 region and avoiding artificial distortions that could compromise anatomical relevance.

Elastic deformations, extreme contrast adjustments, and severe brightness modifications were intentionally excluded to prevent distortion of subtle radiographic patterns related to calcification detection.

Data augmentation was applied exclusively to the training subset. Validation and test images were maintained without any transformations, ensuring an unbiased assessment of the model's generalization performance on unaltered real-world data.

4.6 MobileNetV2 Architecture and Model Configuration

A convolutional neural network based on the MobileNetV2 architecture was adopted as the backbone of the proposed framework for automatic detection of carotid atheromas in panoramic dental radiographs. MobileNetV2 was selected due to its favorable balance between classification performance and computational efficiency, which is particularly suitable for applications involving heterogeneous datasets and real-world clinical screening environments with limited hardware resources.

MobileNetV2 is built upon depthwise separable convolutions combined with inverted residual blocks and linear bottlenecks, significantly reducing the number of trainable parameters and floating-point operations compared with traditional CNN architectures, while maintaining strong representational capacity for visual pattern recognition tasks.

4.6.1 Transfer Learning Initialization

The adopted MobileNetV2 model was initialized with weights pretrained on the ImageNet dataset. Transfer learning enables the reuse of visual features learned from large-scale datasets—including edge detection, texture representation, contrast sensitivity, and basic shape encoding—which remain relevant across different image domains, including medical imaging.

This strategy is particularly advantageous in scenarios involving relatively small annotated datasets, as it:

- Accelerates training convergence;
- Improves generalization capacity;
- Reduces susceptibility to overfitting;

- Enhances stability of the optimization process.

All layers of the pretrained backbone were initially retained, and training was directed toward adapting these representations to the specific characteristics of radiographic patterns associated with carotid atheromas.

4.6.2 Classifier Head Design

To adapt the pretrained backbone to the binary classification task of carotid atheroma detection, the original ImageNet classification layers were removed and replaced with a new task-specific classifier head.

The classifier head consisted of:

- A Global Average Pooling (GAP) layer applied to the final convolutional feature maps, reducing dimensionality while preserving semantic feature representation;
- A fully connected (dense) layer for feature combination;
- A Dropout layer for regularization, aiming to mitigate overfitting;
- A final dense layer with sigmoid activation, generating a probability score representing the likelihood of the presence of carotid calcifications (positive class).

This configuration enables efficient end-to-end training while maintaining compact model complexity compatible with clinical deployment environments.

4.6.3 Input Configuration

All ROIs were resized and formatted to match the MobileNetV2 input requirements. Because the pretrained architecture expects three-channel RGB images and the panoramic radiographs were originally grayscale, channel replication was employed (as described in Section 4.5), producing three identical intensity channels per ROI. This approach ensures full compatibility with the pretrained backbone while preserving the structural integrity of the radiographic information.

4.6.4 Fine-Tuning Strategy

Network training was conducted using a two-stage optimization strategy.

Stage 1 — Warm-up Phase

Initially, all convolutional layers of the MobileNetV2 backbone were frozen, and only the newly added classifier head was trained. This phase aimed to adapt the decision layer parameters to the specific task of carotid atheroma detection while preserving the pretrained convolutional feature representations derived from ImageNet.

Stage 2 — Fine-Tuning Phase

After stabilization of the initial training phase, partial unfreezing of the backbone was performed, allowing fine-tuning of approximately **70% of the upper layers** of the MobileNetV2 network. This approach enabled progressive specialization of deeper convolutional filters toward the identification of radiographic patterns associated with carotid calcifications, while maintaining the generalized low-level visual features encoded by the initial layers.

During this stage, a reduced learning rate was employed to minimize abrupt weight updates and ensure stable convergence. Training regularization and control mechanisms—including Early Stopping and adaptive learning-rate reduction using ReduceLROnPlateau—were applied to prevent overfitting and enhance model stability, as detailed in the following section.

This two-stage training strategy, widely adopted in transfer learning frameworks, enabled a balanced trade-off between task-specific adaptation and preservation of generalized feature extraction, resulting in stable convergence and satisfactory discriminative performance despite the limited dataset size and inherent heterogeneity of panoramic radiographs. This configuration supports the practical feasibility of the proposed deep learning model for clinical opportunistic screening applications.

4.7 Training Strategy

The training process was carefully structured to ensure stable convergence, mitigation of overfitting, and robust evaluation of the neural network's generalization capability. All computational procedures were implemented in Python using the TensorFlow/Keras framework, within a computing environment equipped with an Intel Core i7 processor and an NVIDIA GeForce RTX 2060 graphics processing unit (6 GB VRAM), enabling efficient training of deep neural networks using high-resolution medical images.

4.7.1 Data Partitioning into Training, Validation, and Test Sets

The final curated dataset was stratified and split into three independent subsets:

- **Training set:** 264 images (157 negative and 107 positive) — approximately 70%;
- **Validation set:** 57 images (34 negative and 23 positive) — approximately 15%;
- **Test set:** 57 images (34 negative and 23 positive) — approximately 15%.

Stratification ensured similar proportions of positive and negative samples across all partitions, preserving the natural class distribution and preventing sampling bias during both training and evaluation phases.

The training subset was used for weight updating during both warm-up and fine-tuning stages. The validation subset served to monitor model generalization on unseen data and was applied for hyperparameter adjustment as well as determination of the optimal early stopping point. The independent test subset was reserved exclusively for the final performance evaluation of the trained model.

This data partitioning strategy ensures appropriate separation between datasets used for learning, parameter tuning, and definitive evaluation, thereby preventing information leakage and avoiding optimistic bias in performance estimation.

4.7.2 Loss Function

Given the binary classification task (presence versus absence of carotid atheromas), the Binary Cross-Entropy (BCE) loss function was adopted and is defined as:

$$L = -[y \cdot \log(p) + (1 - y) \cdot \log(1 - p)],$$

where y represents the true class label (0 or 1), and p denotes the predicted probability of the positive class output by the network.

This loss function is widely recommended for probabilistic binary classification problems due to its effectiveness in directly measuring the discrepancy between predicted and true class probabilities while providing stable gradients for network optimization.

To mitigate the influence of the mild class imbalance observed in the training dataset, the loss function was weighted using class weights, as described in Section 4.7.4.

4.7.3 Optimizer

Model training was performed using the Adam (Adaptive Moment Estimation) optimizer, which is broadly recognized for its robustness and efficiency in deep learning applications. Adam combines the benefits of momentum-based optimization with adaptive learning rates for individual parameters, promoting faster and more stable convergence.

During the initial warm-up phase, a learning rate of 1×10^{-3} was employed to facilitate rapid adaptation of the newly added classifier head layers.

In the subsequent fine-tuning phase, the learning rate was reduced to 1×10^{-5}, enabling gradual refinement of the weights in the upper convolutional layers of the pretrained backbone without compromising previously learned feature representations.

As an additional training control strategy, ReduceLROnPlateau was applied with a reduction factor of 0.5 and a patience of three epochs, automatically decreasing the learning rate when validation performance plateaued, down to a minimum limit of 1×10^{-6}.

4.7.4 Class Weighting

Due to the asymmetric distribution of classes in the training dataset, with a predominance of negative samples, class weighting was implemented during loss optimization to balance the contribution of each class.

Class weights were calculated inversely proportional to class frequencies in the training set and resulted in the following values:

- **Negative class (0): 0.84**
- **Positive class (1): 1.23**

These weighting factors increased the penalty associated with misclassification of positive samples (atheroma cases), thereby promoting higher model sensitivity and reducing bias toward the majority class.

The application of class weighting, combined with the real-time data augmentation strategy detailed in Section 4.5.3, contributed to balancing the learning process without the need for explicit oversampling or synthetic replication of training examples, while preserving the natural variability of the dataset.

Together, these carefully designed training procedures enabled a stable learning process and effective generalization across heterogeneous panoramic radiographs. This strategy supports the practical applicability of the proposed deep learning model for opportunistic clinical screening, reinforcing its potential role as a scalable auxiliary tool for early detection of carotid atherosclerotic disease in routine dental settings.

4.8 Performance Evaluation

Evaluation of the proposed model was conducted with the objective of comprehensively and clinically assessing its discriminative ability for identifying carotid atheromas in panoramic dental radiographs. For this purpose, standard binary classification metrics (Esteva et al., 2021), graphical analyses using ROC and precision–recall curves (Saito; Rehmsmeier, 2015), and confusion matrices were employed, all derived exclusively from predictions on the independent test set (Jin et al., 2023).

All reported metrics were calculated solely on the test partition, which was not used at any stage of training, fine-tuning, or hyperparameter optimization, thereby ensuring an unbiased estimation of model performance on previously unseen data.

4.8.1 Performance Metrics

The following metrics were used to characterize model performance. In all expressions, **TP** denotes true positives, **TN** true negatives, **FP** false positives, and **FN** false negatives.

Accuracy (ACC) — global proportion of correct classifications:

$$ACC = \frac{TP + TN}{TP + TN + FP + FN}$$

Sensitivity (SEN, Recall) — ability of the model to correctly identify positive cases (atheromas):

$$SEN = \frac{TP}{TP + FN}$$

Specificity (SPE) — ability of the model to correctly classify negative cases:

$$SPE = \frac{TN}{TN + FP}$$

Precision (PRE, Positive Predictive Value – PPV) — proportion of true positives among all positive predictions:

$$PRE = \frac{TP}{TP + FP}$$

F1-score — harmonic mean of precision and sensitivity, particularly useful for imbalanced datasets:

$$F1 = \frac{2 \cdot PRE \cdot SEN}{PRE + SEN}$$

Area Under the ROC Curve (AUC) — evaluates the global discriminative capacity of the classifier across all possible decision thresholds.

Area Under the Precision–Recall Curve (AUPRC) — a metric especially relevant for moderately imbalanced datasets, as it reflects the trade-off between sensitivity and precision in positive-case detection.

4.8.2 ROC Curves and Confusion Matrices

The Receiver Operating Characteristic (ROC) curve was constructed to analyze the discriminative behavior of the classifier across different decision thresholds, plotting sensitivity (true positive rate) against the false positive rate (1 − specificity).

Similarly, precision–recall curves were generated to evaluate the relationship between precision and sensitivity, which is particularly informative under class imbalance conditions.

Confusion matrices were used to present classification results at selected operating points, providing a direct numerical breakdown of TP, TN, FP, and FN values. These matrices allowed straightforward visualization of error patterns and supported clinical interpretation of the model's classification behavior.

Together, these tools enabled comprehensive evaluation of both the overall diagnostic performance and the operational trade-offs relevant to practical screening applications.

4.8.3 Validation Procedures and Threshold Analysis

Model selection during training was guided by monitoring the validation ROC area (val-AUC). The checkpoint associated with the highest validation AUC was preserved and applied in the final evaluation on the independent test set.

The definitive test analysis was conducted using two principal operational criteria:

Youden's Index Threshold

The optimal decision threshold was identified using the Youden index:

$$J = SEN + SPE - 1$$

This criterion selects the threshold that maximizes the combined balance between sensitivity and specificity and is appropriate for evaluating overall classifier performance.

Sensitivity-Target Threshold (SEN ≈ 0.80)

A second operational threshold was defined to achieve a target sensitivity near **80%**, simulating a **screening-oriented clinical scenario** in which minimizing false negatives is prioritized, even at the expense of a moderate reduction in specificity.

This approach enabled performance evaluation under two distinct clinical-use conditions:

- **Balanced operation** (Youden J optimization) (Youden, 1950);
- **Screening-oriented operation**, focused on enhanced detection capability with predefined sensitivity targets.

Statistical Uncertainty Estimation

To quantify the statistical uncertainty associated with performance estimates, a **bootstrap resampling procedure with 1,000 iterations** was applied to the test dataset, generating **95% confidence intervals (CIs)** for the primary performance metrics (ACC, SEN, SPE, PPV, AUC, and AUPRC).

At each iteration, samples were drawn with replacement from the test set, and metrics were recomputed, allowing estimation of empirical distributions and variability of the results. This procedure ensured robust performance reporting and increased confidence in the stability and reproducibility of the experimental findings.

4.9 Bootstrap Analysis and Decision Threshold Optimization

The final stage of the methodological evaluation aimed to quantify the statistical uncertainty associated with the model's performance metrics and to define clinically meaningful decision thresholds for classifier operation in vascular risk screening scenarios. For this purpose, nonparametric bootstrap resampling techniques and systematic threshold optimization procedures were employed.

4.9.1 Bootstrap Resampling

Given the relatively limited size of the independent test set and the need to obtain robust performance estimates, a nonparametric bootstrap method was applied. This procedure consisted of generating 1,000 resampled datasets with replacement from the original test set.

For each resampled subset, performance metrics—including Accuracy (ACC), Sensitivity (SEN), Specificity (SPE), Precision (PPV), Area Under the ROC Curve (AUC), and Area Under the Precision–Recall Curve (AUPRC)—were recomputed. The resulting metric values across iterations were used to construct empirical distributions for each measure.

Bootstrap resampling is particularly appropriate in scenarios where normal distribution assumptions cannot be safely made and when small or moderate sample sizes limit the applicability of traditional parametric inference methods. This approach therefore provides a robust estimation of statistical variability for the reported results.

4.9.2 Confidence Interval Estimation

Based on the empirical metric distributions generated by bootstrap resampling, 95% confidence intervals (95% CIs) were estimated for all evaluated performance measures using the percentile method, corresponding to the 2.5% and 97.5% quantiles of the bootstrap distributions.

The presentation of confidence intervals enabled:

- Assessment of classifier performance stability;
- Quantification of statistical variability associated with point estimates;
- Reduction of interpretational bias derived from isolated performance values;
- Support for methodological comparison with previously published studies.

4.9.3 Selection of Operational Decision Thresholds

For each ROI, the MobileNetV2 classifier outputs a continuous class probability $p \in [0,1]$representing the likelihood of carotid atheroma presence. Conversion of these probabilities into binary class assignments required the definition of operational decision thresholds. Two threshold selection strategies were applied:

a) Youden Index Optimization

The first decision threshold was determined based on the Youden index (J) (Youden, 1950), defined by:

$$J = SEN + SPE - 1$$

The selected threshold corresponds to the point along the ROC curve that maximizes the combined balance between sensitivity and specificity, representing the classifier's optimal global operating point. This criterion was adopted to report the model's standard balanced performance in the simultaneous detection of positive and negative cases.

b) Sensitivity-Target Threshold

Considering the clinical context of opportunistic screening, where minimizing false-negative results is prioritized, a second decision threshold was selected to achieve a target sensitivity of approximately 80% (SEN ≈ 0.80).

This operating point reflects a screening-oriented clinical scenario, in which enhanced detection of potentially affected individuals is favored, even at the cost of a controlled increase in false-positive rates. Such a strategy is consistent with the goals of early triage and referral for specialized vascular assessment.

The adoption of both thresholds allowed evaluation of the classifier under two clinically relevant operational configurations:

- **Balanced mode** — threshold optimized using the Youden index;
- **Screening mode** — sensitivity-prioritized operation with a predefined detection target.

In both scenarios, all performance metrics were independently computed on the test set, accompanied by their respective bootstrap-derived confidence intervals.

4.9.4 Comparison with Previous Literature

The results obtained at both operational thresholds were qualitatively compared with data reported in prior studies addressing automatic detection of carotid atheromas on panoramic dental radiographs.

Under balanced operating conditions, the proposed method demonstrated performance comparable to or superior to most previously published approaches, particularly with respect to achieving a simultaneous combination of high sensitivity and specificity—an outcome frequently identified as challenging in earlier investigations. Under screening-oriented conditions, the classifier maintained high specificity even while enforcing a minimum sensitivity requirement, consistent with protocols advocated in clinical risk screening strategies.

These analyses reinforce the methodological suitability of the selected thresholding approaches and highlight the operational flexibility of the proposed system, enabling adaptation to specific clinical objectives and healthcare resource constraints.

4.10 Code Availability / Reproducibility

The source code developed for this study, including preprocessing routines, model training, and evaluation scripts, is publicly available on GitHub at:

- *https://github.com/thais-ufu/carotid-atheroma-screening-panoramic-xray*

The repository contains the complete implementation used in the experiments reported in this dissertation and allows full reproducibility of the results.

All scripts were implemented in Python using TensorFlow/Keras.

5 RESULTS

This section presents the experimental results obtained with the MobileNetV2-based convolutional neural network for detecting carotid atheromas in panoramic dental radiographs. Results are organized as follows: (i) training and validation behavior, (ii) confusion matrix and classification metrics on the independent test set, and (iii) comparative analysis with previous studies. Unless otherwise specified, model selection was based on the highest validation AUC via checkpointing. The dataset partitions followed the stratified protocol described in Section 4. The operating point on the test set is reported both at the Youden J threshold (Youden, 1950) and at a target sensitivity of ≈0.80. The independent test set is detailed in Section 3.2.

5.1. Model Performance

The learning dynamics of the MobileNetV2 model were examined to assess convergence and generalization. Training followed two phases: (i) head training with the backbone frozen and (ii) fine-tuning with partial unfreezing of the last convolutional blocks.

Figure 7**Erro! Fonte de referência não encontrada.** and Figure 8 summarize the evolution of training/validation accuracy and loss. Validation curves closely tracked training curves throughout both phases, indicating that augmentation, class weighting, and regularization successfully mitigated overfitting despite the limited dataset.

Training and validation AUC across epochs for MobileNetV2. Validation AUC increased steadily from ~0.60 at the beginning of phase 1 to 0.98 at the best epoch in phase 2, while training AUC approached ~0.99, indicating stable convergence with good generalization.

Figure 7 Training and validation AUC across epochs for MobileNetV2. Validation AUC increased steadily from ~0.60 at the beginning of phase 1 to 0.98 at the best epoch in phase 2, while training AUC approached ~0.99, indicating stable convergence with good generalization.

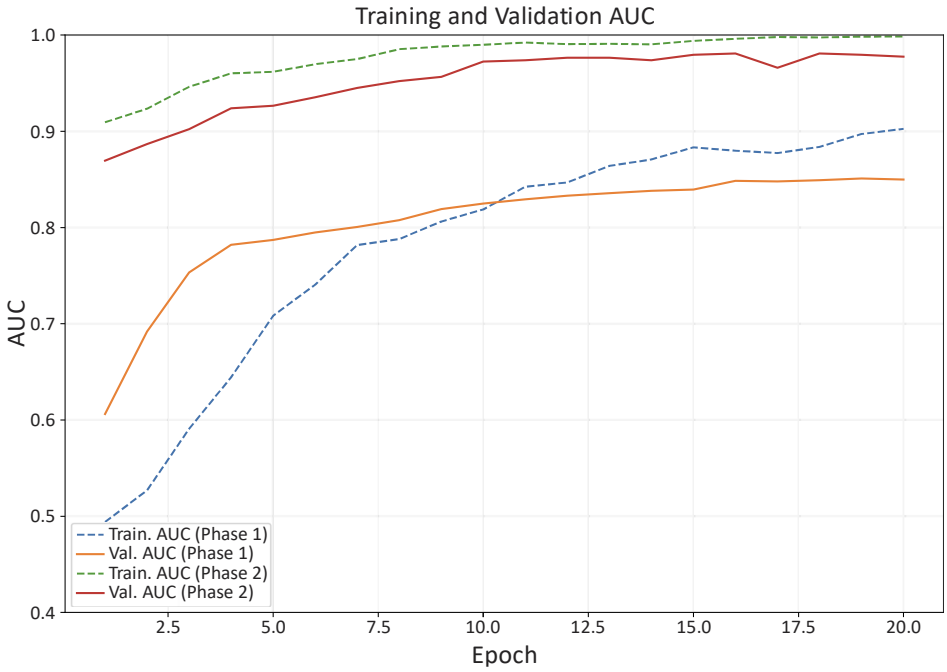

Source: Author's own elaboration based on the experimental results.

Figure 8 Training and validation loss across epochs. Training loss decreased monotonically through both phases, and validation loss, although noisier after fine-tuning, reached its minimum in phase 2 (≈0.24–0.30), consistent with the AUC peak and the selected checkpoint.

Source: Author's own elaboration based on the experimental results.

Table 1 reports representative checkpoints from the logs, emphasizing the primary selection metric (validation AUC) together with accuracy and loss. These results indicate stable convergence, with fine-tuning of the upper convolutional blocks yielding the highest discrimination on validation (AUC = 0.95), which guided model selection for subsequent test-set evaluation.

Table 1. Training/validation performance at representative epochs (primary metric = val AUC).

Phase	Epoch	Train ACC	Val ACC	Val AUC	Train Loss	Val Loss
1 - Frozen	1	0.58	0.53	0.58	0.77	0.70
	6	0.63	0.74	0.77	0.67	0.63
	13	0.67	0.60	0.81	0.62	0.68
	19	0.72	0.60	0.84	0.57	0.70
2 - Fine-tuning	1	0.75	0.47	0.84	0.51	0.79
	3	0.81	0.61	0.85	0.49	0.71
	10	0.80	0.54	0.86	0.48	0.84
	13*	**0.97**	**0.70**	**0.95**	**0.14**	**0.73**

* Best validation AUC (checkpoint used for evaluation).

Source: Author's own elaboration based on the experimental results.

5.2. Confusion Matrix and Metrics

The performance of the proposed MobileNetV2 classifier was evaluated on an independent test set comprising 57 radiographs (34 negatives, 23 positives). Results are reported at two operating points: (i) the Youden J threshold estimated on the validation split and (ii) a target-sensitivity operating point ($SEN \approx 0.80$). The corresponding metrics are summarized in Table 2.

Table 2. Classification metrics on the independent test set.

Metric	Youden J (thr ≈ 0.84)	SEN ≈ 0.80 (thr ≈ 0.98)
Accuracy (ACC)	94.7%	91.2%
Precision / Positive Predictive Value (PPV)	91.7%	95.0%
Sensitivity (SEN)	95.7%	82.6%
Specificity (SPE)	94.1%	97.1%
F1-score (F1)	93.6%	88.3%
Area Under the ROC Curve (AUC)	0.963	0.968
Area Under the Precision–Recall Curve (AUPRC)	0.968	0.970

Source: Author's own elaboration based on the experimental results.

At the Youden operating point, the model achieved 94.7% accuracy with balanced and high sensitivity (95.7%) and specificity (94.1%), and strong discrimination (AUC 0.96, AUPRC 0.97). This setting prioritizes overall diagnostic balance and minimizes total classification error. At the target-sensitivity operating point ($SEN \approx 0.80$), the model maintained 91.2% accuracy with 97.1% specificity and 95.0% precision, which may be preferable when aiming to reduce false positives while still preserving high sensitivity. These results indicate robust generalization on the independent test set. Given the dataset size, confidence intervals were computed via bootstrapping in the evaluation scripts to account for sampling variability, and the trends remained consistent across thresholds. The confusion matrices are shown in Figure 9.

Figure 9 - Confusion matrices on the test set.

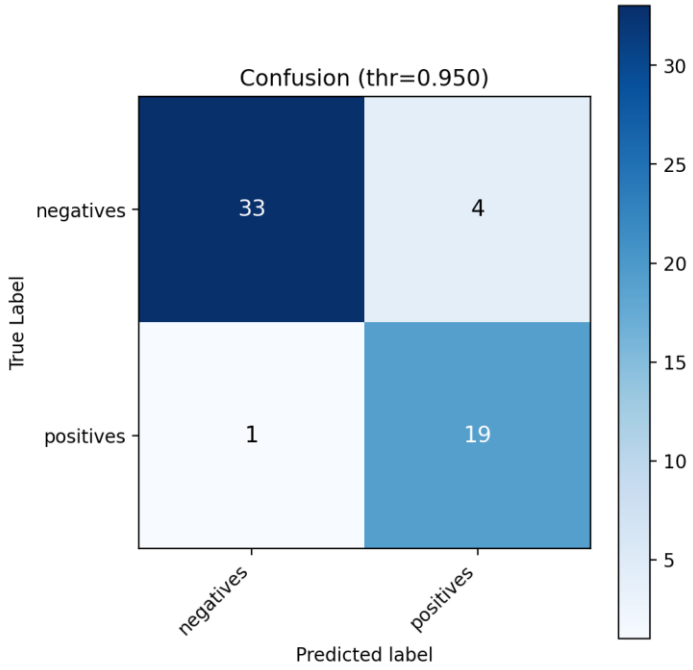

Source: Author's own elaboration based on the experimental results.

5.3. Comparative Analysis

To contextualize the results, the performance of the proposed MobileNetV2 classifier was contrasted with prior studies on automated detection of carotid atheromas in panoramic radiographs. Although datasets, preprocessing pipelines, and network designs vary widely, the literature commonly reports accuracies in the 80–92% range, with sensitivity and specificity often imbalanced, as outlined in (Janiszewska-Olszowska et al., 2022).

Under the Youden J operating point on the independent test set, the present model achieved *ACC* = 94.7%, *SEN* = 95.7% *SPE* = 94.1%, and AUC = 0.96. These values are comparable to, and in some respects exceed, those reported by approaches based on handcrafted features and classical machine-learning classifiers (typically ~75–85% accuracy) (Ikeda et al., 2024; Kats et al., 2019; Vinayahalingam et al., 2024). Recent CNN-based methods have surpassed 85% accuracy (Kats et al., 2019; Vinayahalingam et al., 2024) but frequently exhibit lower specificity, suggesting a propensity for false positives. In contrast, the current results show balanced performance across classes, indicating effective control of bias between positive and negative predictions.

The combination of MobileNetV2 with ROI-standardized preprocessing, cost-sensitive training (class weights), and conservative data augmentation appears to contribute to this balance. While the modest test-set size warrants cautious interpretation and precludes definitive statistical comparison with larger cohorts, the findings align with the state of the art and reinforce the feasibility of deep learning for opportunistic identification of carotid atheromas on panoramic radiographs.

Figure 10 presents a quantitative comparison between the proposed approach and prior studies (Ikeda et al., 2024; Kats et al., 2019; Vinayahalingam et al., 2024). The MobileNetV2 model outperformed earlier CNN and detector-based methods, achieving a higher overall balance between sensitivity and specificity. These comparisons motivate a broader discussion of clinical applicability, methodological limitations, and avenues for future research, addressed in the next section.

Figure 10. Comparative performance of carotid atheroma detection models reported in previous studies versus the proposed MobileNetV2 model. The bars represent the ACC, SEN, SPE, and AUC, where available. The proposed approach achieved balanced and superior performance across all metrics while maintaining computational efficiency.

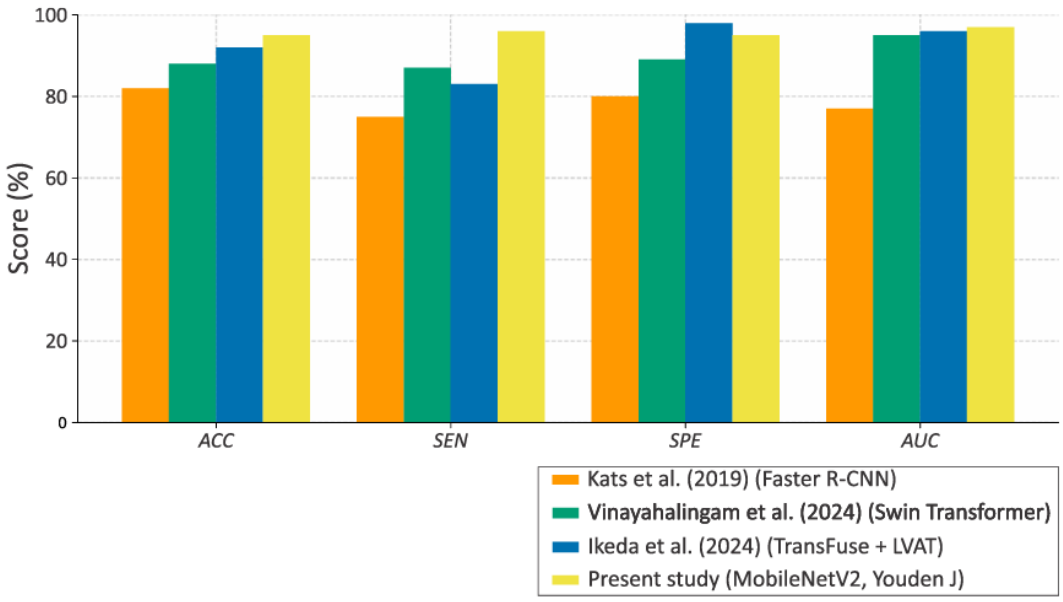

Source: Data compiled by the authors based on results reported in (Kats et al., 2019), (Vinayahalingam et al., 2024), (Ikeda et al., 2024) and the present study.

5.4. Qualitative Error Assessment

To provide qualitative insight into the model's behavior and to contextualize the quantitative results, Figure 11 presents representative examples from the independent test set, including true positive (TP), true negative (TN), false positive (FP), and false negative (FN) cases. This qualitative analysis complements the numerical performance metrics by illustrating both successful detections and clinically relevant failure modes encountered during carotid atheroma screening on panoramic radiographs.

Figure 11. Representative success and failure cases on the independent test set: (**A**) True positive case showing a calcified carotid atheroma correctly identified by the model. (**B**) True negative case without calcifications in the carotid bifurcation region. (**C**) False positive case caused by an overlapping elongated anatomical structure within the carotid ROI, mimicking a calcified lesion. (**D**) False negative case illustrating a subtle and anatomically ambiguous calcification near the mandibular border, with differential diagnosis including phleboliths and carotid artery calcifications. Circles and arrows are manually added for visualization purposes only.

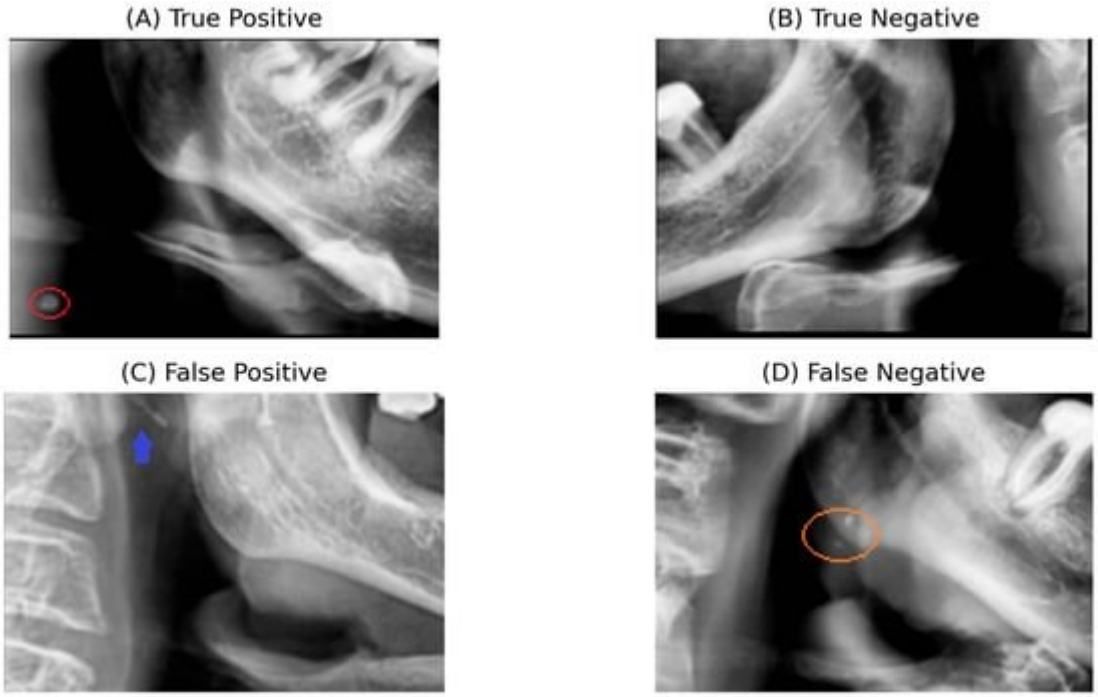

Source: Adapted from the public dataset by Mureșanu; Hedeșiu; Iacob, (2025). ROIs were extracted and visual annotations (circles and arrows) were manually added by the author for illustrative purposes.

False positive predictions were predominantly observed in cases involving uncommon radiographic patterns in which elongated or overlapping structures appeared within the carotid ROI. Although a definitive causal explanation cannot be established based solely on qualitative inspection, these findings are consistent with well-

documented anatomical configurations that may mimic calcified atheromatous plaques on panoramic radiographs. Minor ROI misalignment or projection variability inherent to panoramic imaging may further contribute to such appearances.

The false negative example represents a radiographically challenging case characterized by a subtle calcification located in close proximity to the mandibular border and adjacent soft-tissue structures. In this scenario, the visual features overlap with common differential diagnoses, including phleboliths and other cervical calcifications, leading to diagnostic ambiguity even for expert readers.

Importantly, these misclassifications should be interpreted as plausible hypotheses grounded in radiographic appearance rather than definitive causal explanations. They reflect intrinsic limitations of panoramic radiography and qualitative assessment, rather than evidence of systematic model bias or instability. From a screening-oriented perspective, prioritizing sensitivity over strict specificity remains a clinically acceptable trade-off, as false-positive findings can be resolved through confirmatory vascular imaging, whereas false-negative outcomes may delay referral of potentially at-risk patients.

6 DISCUSSION

6.1 Synthesis of Principal Findings

This section synthesizes the main experimental findings obtained in the evaluation of the proposed deep learning framework. This study demonstrated that the proposed method is effective for the detection of carotid atheromas on panoramic dental radiographs, confirming the feasibility of using lightweight convolutional neural networks as tools for opportunistic vascular screening in dental settings. The combination of anatomically standardized ROIs in the C3–C4 region, conservative preprocessing, and a two-stage MobileNetV2 fine-tuning strategy resulted in consistent discrimination between ROIs with and without radiographic signs of carotid calcification. These results align with recent trends reported in the literature emphasizing the applicability of efficient deep learning architectures for scalable screening solutions in medical and dental imaging (Aggarwal et al., 2021; Vinayahalingam et al., 2024).

From a clinical perspective, the achieved performance can be considered satisfactory for screening-oriented applications. The model was capable of maintaining a balanced trade-off between sensitivity and specificity, ensuring that a high proportion of potentially affected individuals could be flagged while preserving an acceptable level of false-positive control. This balance is particularly relevant in opportunistic screening contexts (Bushnell et al., 2024), where both the avoidance of missed cases and the prevention of excessive unnecessary referrals must be taken into account.

A notable strength of the method lies in its ability to operate under different clinical operating modes. When optimized for overall performance using a balanced decision threshold, the classifier achieved simultaneous optimization of detection and exclusion capabilities. Conversely, when configured for screening-oriented operation, with an emphasis on enhanced sensitivity, the system maintained adequate specificity, demonstrating flexibility in adapting to different healthcare workflows and resource-availability scenarios (Alsakar et al., 2024; Anwar et al., 2018). This dual operational capability supports the clinical applicability of the approach beyond a single fixed threshold setting.

Furthermore, the model exhibited good generalization performance, despite the use of heterogeneous image sources obtained from different institutions and imaging

devices. This robustness suggests that the standardized ROI pipeline successfully minimized device- and protocol-related variability, allowing the classifier to focus on anatomically relevant patterns rather than dataset-specific noise (Aggarwal et al., 2021; Vinayahalingam et al., 2024).

Finally, the application of bootstrap resampling analysis confirmed the stability of the obtained performance metrics, providing statistical support for the consistency of the model's predictions across repeated test set perturbations. The narrow dispersion of bootstrap estimates reinforces the reliability of the reported discrimination capacity and reduces uncertainty associated with the relatively modest test set size. Taken together, these findings validate the methodological choices adopted in this study and support the conclusion that deep learning–based analysis of panoramic radiographs represents a promising strategy for opportunistic identification of carotid atheromas in routine dental imaging. These observations provide the basis for the more detailed methodological and comparative discussion presented in the following sections.

6.2 Methodological Determinants of Performance

The strong performance observed in this study can be attributed to a set of complementary methodological choices that collectively enhanced model robustness, learning stability, and clinical applicability.

A central determinant was the standardization of regions of interest (ROIs) focused on the anatomical level of C3–C4. By restricting the input to the carotid bifurcation region, visual noise from unrelated dental and maxillofacial structures was substantially reduced. This targeted cropping strategy constrained the learning process to anatomically meaningful features and minimized spatial variability between samples, thereby promoting training stability and more efficient feature extraction. In contrast to approaches that rely on full panoramic images, the ROI-based pipeline effectively guided the network's attention toward the expected location of calcifications, reducing the risk of spurious correlations and background-driven misclassifications. This anatomical ROI strategy is consistent with trends reported in recent CNN-based vascular screening studies, which emphasize spatial constraint as a key factor for improving diagnostic specificity and model robustness (Prados-Privado et al., 2022; Vinayahalingam et al., 2024).

The adoption of MobileNetV2 as the backbone architecture was another key methodological factor. MobileNetV2 provides a favorable balance between representational capacity and computational efficiency (Sandler et al., 2019) through depthwise separable convolutions and inverted residual blocks. This design enables the network to capture fine radiographic patterns while maintaining a relatively low parameter count. Such efficiency is particularly advantageous when training on moderate-sized datasets, where larger architectures may be prone to overfitting or require aggressive regularization. Furthermore, its lightweight nature supports the clinical feasibility of real-time or near–real-time deployment within resource-limited dental environments, aligning with current recommendations for scalable artificial intelligence solutions in routine diagnostic workflows.

The application of transfer learning further contributed to performance optimization. Initializing the network with weights pre-trained on large-scale image datasets allowed the early layers to retain generic feature detectors for edges, contours, and textures, which remain relevant for radiographic interpretation. This strategy promoted faster and more stable convergence compared to training from scratch and reduced dependency on large annotated datasets. Given the moderate sample size available for this study, transfer learning was essential for achieving high discriminative performance while minimizing the risk of unstable training dynamics, a benefit widely described in the contemporary medical imaging literature (Aggarwal et al., 2021).

Another relevant methodological component was the use of class weighting during optimization to compensate for the mild imbalance between negative and positive samples. By assigning higher loss penalties to misclassified positive cases, the training process was biased toward enhanced detection capability, thus contributing directly to the observed improvement in sensitivity without substantially compromising specificity. This approach aligns well with the priorities of screening applications, where minimizing false negatives is clinically critical and has been recommended in prior studies focusing on AI-assisted detection of high-risk conditions (Vinayahalingam et al., 2024).

Finally, the implementation of realistic on-the-fly data augmentation played a fundamental role in reducing overfitting and improving generalization (Aggarwal et al., 2021). Mild geometric transformations — such as horizontal flipping, small rotations, translations, and zoom variations — simulated natural acquisition variability associated

with patient positioning and imaging geometry. Importantly, these augmentations were conservative, preserving anatomical integrity of the carotid region while expanding visual diversity of training samples. This strategy increased the effective dataset size and exposed the network to a broader distribution of plausible inputs, thereby fostering greater robustness to unseen variations and contributing to the consistent test-set performance observed in this study.

Collectively, these methodological determinants — anatomical ROI standardization, efficient network architecture, transfer learning initialization, balanced loss optimization through class weighting, and realistic data augmentation — formed a synergistic framework. Rather than relying on any isolated component, the favorable performance emerged from the interaction between design choices across the entire pipeline. This integrated methodological approach helps explain the observed stability, diagnostic balance, and scalability of the proposed system within realistic screening contexts.

6.3 Comparison with Literature

The findings obtained in this study are broadly consistent with previous investigations that explored the use of deep learning techniques for detecting carotid artery calcifications and related soft-tissue findings on panoramic dental radiographs. Over the last decade, multiple studies have demonstrated that convolutional neural networks are capable of identifying radiographically visible carotid calcifications with clinically meaningful accuracy, reinforcing the potential role of dental imaging as an opportunistic screening tool for vascular risk assessment (Amitay et al., 2023; Arzani et al., 2025; Vinayahalingam et al., 2024).

From a performance standpoint, the discriminative ability achieved by the proposed model, as reflected by high values of area under the ROC and Precision–Recall curves, lies within the upper range reported in the literature. Systematic reviews and meta-analyses have indicated that CNN-based approaches applied to medical and dental imaging commonly reach AUC values above 0.85, with more recent studies approaching or exceeding 0.90 when carefully curated datasets and modern architectures are employed (Aggarwal et al., 2021; Prados-Privado et al., 2022). In this context, the results obtained in the present work are aligned with state-of-the-art performance levels reported for carotid atheroma detection on panoramic radiographs.

Despite this overall agreement, relevant differences emerge when examining how prior studies balance sensitivity and specificity. Several published approaches emphasize high detection sensitivity, often at the cost of elevated false-positive rates, which may lead to unnecessary medical referrals in a screening scenario (Kats et al., 2019; Song et al., 2022). Conversely, other studies report high specificity but reduced sensitivity, increasing the risk of missed carotid atheromas and limiting their utility for preventive screening (Vinayahalingam et al., 2024). In contrast, the present study demonstrated a more symmetrical sensitivity–specificity trade-off, suggesting a more favorable equilibrium between identifying true positive cases and avoiding excessive false alarms. This balanced behavior is particularly relevant in opportunistic screening contexts, where both false negatives and false positives carry important clinical and economic implications.

The methodological design adopted in this work appears to play a central role in achieving this equilibrium. Specifically, the use of standardized anatomical ROI cropping centered on the C3–C4 region reduces background variability and ensures consistent representation of the carotid bifurcation across samples, a limitation frequently acknowledged in earlier studies relying on loosely defined or heterogeneous regions of interest (Felipe et al., 2020; Manta et al., 2024). In addition, the application of class-weighted optimization mitigates the effects of class imbalance, a recurrent challenge in carotid calcification datasets due to the relatively low prevalence of positive findings in routine dental populations (Arzani et al., 2025).

Methodological limitations commonly reported in the literature further help explain disparities in performance across studies. A substantial portion of prior work relied on small or institution-specific datasets, often acquired under homogeneous imaging protocols. While such conditions may inflate internal validation results, they also increase susceptibility to overfitting and reduce generalizability to real-world clinical environments (Anwar et al., 2018; Zhou et al., 2021). Additionally, some investigations employed high-capacity detection networks or transformer-based architectures, which, although expressive, require large annotated datasets to achieve stable generalization and often impose significant computational demands, limiting their feasibility for deployment in routine dental practice.

Another important distinction concerns evaluation methodology. Several earlier studies reported point estimates of accuracy, sensitivity, or specificity without providing

confidence intervals or resampling-based uncertainty analysis. The absence of uncertainty quantification complicates cross-study comparisons and may obscure the robustness of reported results, particularly when test sets are small. In contrast, the present study incorporated bootstrap resampling procedures to estimate confidence intervals for key performance metrics, enabling a more rigorous assessment of model stability and predictive reliability.

Overall, the comparison with existing literature indicates that the proposed approach matches or surpasses previous methods in key diagnostic metrics while offering advantages in operational balance, statistical characterization, and computational efficiency. Rather than prioritizing architectural complexity, this work demonstrates that a carefully designed pipeline—combining anatomical ROI standardization, moderate-capacity convolutional backbones, transfer learning, loss weighting, and conservative data augmentation—can achieve state-of-the-art diagnostic performance with enhanced robustness. Importantly, this methodological strategy supports practical feasibility for deployment in real-world dental workflows, strengthening the translational potential of deep learning–based opportunistic screening for carotid atheromas.

6.4 Clinical Implications

The findings of this study indicate that the proposed method is not only technically valid but also clinically feasible, supporting its potential use as an assistive screening tool in real-world dental practice. The reliance on panoramic radiographs—an examination already routinely acquired for dental diagnosis—enables implementation of the method without introducing additional imaging procedures, additional radiation exposure, or significant workflow burdens. This characteristic positions the approach as an opportunistic screening strategy that can be seamlessly incorporated into standard dental care.

The most appropriate clinical scenario for application is the dental outpatient setting, where panoramic radiography is commonly obtained for baseline assessment, orthodontic planning, implant evaluation, or surgical follow-up. Previous studies have emphasized that panoramic images frequently include the anatomical region of the carotid bifurcation and therefore offer a unique opportunity for incidental detection of carotid calcifications during routine dental care (Arzani et al., 2025; Prados-Privado et

al., 2022; Schiroli et al., 2025). In this context, the proposed algorithm could operate in the background during routine image review, automatically analyzing predefined ROIs and providing a probability score for the presence of carotid calcifications. When this probability exceeds a predefined decision threshold—particularly one configured to prioritize sensitivity—the system could flag the examination for closer manual review or suggest referral for medical evaluation.

Such a triage-oriented workflow does not convert dental imaging into a definitive diagnostic pathway for vascular disease. Panoramic radiography cannot replace established vascular imaging modalities, such as Doppler ultrasonography, computed tomography angiography, or magnetic resonance angiography, which remain the gold standards for the diagnosis and characterization of carotid atherosclerosis and stenosis (Daolio et al., 2024; David et al., 2024; Saxena; Ng; Lim, 2019). Instead, the proposed method should be understood as an early warning mechanism, capable of identifying individuals who may benefit from further investigation using appropriate medical imaging techniques.

By facilitating early referral of at-risk individuals, the proposed tool may contribute to reducing delays in diagnosis and enabling preventive interventions before the occurrence of cerebrovascular events. This potential impact is particularly relevant given the well-documented burden of stroke on public health systems worldwide and the recognized importance of early identification of vascular risk factors (Abissegue et al., 2024; Bushnell et al., 2024; World Health Organization, 2011). Opportunistic screening strategies that leverage routinely acquired examinations have been highlighted as promising approaches to improve population-level prevention without increasing healthcare costs or patient burden.

A key aspect underpinning the practicality of this approach is the selection of the MobileNetV2 architecture. Its lightweight design, characterized by reduced computational complexity and low memory requirements, enables efficient execution on conventional clinical workstations without the need for dedicated graphics processing units or cloud-based infrastructures. Prior studies have demonstrated that MobileNet-based architectures can achieve competitive diagnostic performance in medical and dental imaging tasks while maintaining computational efficiency, supporting their suitability for real-world clinical deployment (Alsakar et al., 2024; Zhou et al., 2021).

Furthermore, the simplicity of implementation—based on standardized ROI extraction, limited preprocessing steps, and a single lightweight classifier—favors scalable deployment in both private dental clinics and public healthcare services, including settings with restricted technical resources. The absence of specialized hardware requirements and minimal disruption to existing workflows enhance the economic and logistical feasibility of integrating the proposed system into everyday clinical routines (Aggarwal et al., 2021; Corbella; Srinivas; Cabitza, 2021).

In summary, the proposed approach demonstrates strong potential as a practical, low-cost, and scalable supportive screening tool within dental practices, particularly in public health contexts where access to specialized vascular imaging is limited. While it does not replace formal vascular diagnostic examinations, it serves as an early detection mechanism capable of bridging dentistry and general medicine, fostering interdisciplinary collaboration and strengthening preventive healthcare strategies.

6.5 Limitations

Despite the encouraging results obtained in this study, several limitations must be acknowledged to ensure proper interpretation of the findings and to guide future research.

First, the size of the dataset remains moderate relative to the demands of deep learning methods. Although the use of transfer learning, class-weighted optimization, conservative data augmentation strategies, and bootstrap-based validation mitigated risks associated with limited data volume, a larger dataset would be desirable to enhance representativeness and further stabilize performance estimates, particularly across less frequent radiographic patterns and rare presentations of calcifications.

Second, the study lacks external validation on fully independent, multi-institutional cohorts beyond the public datasets used for model development and testing. Although the combined datasets included images acquired from different institutions and imaging devices, all testing was performed within the same overall data distribution framework used for training and hyperparameter selection. External validation across geographically and operationally distinct clinical environments is widely recognized as essential to confirm real-world generalizability of artificial intelligence models in medical imaging and to assess potential domain-shift effects (Aggarwal et al., 2021; Zhou et al., 2021). Such effects may arise from regional practice variations, differences

in patient demographics, or unfamiliar acquisition protocols, and are particularly relevant in low-resource or public-health settings.

Third, the research design was retrospective in nature, relying exclusively on previously collected radiographic data rather than prospective clinical workflows. As a result, real-time operational aspects—such as integration into dental imaging software, effects on consultation dynamics, clinician acceptance, and adherence to referral recommendations—were not evaluated. Prior studies in medical artificial intelligence emphasize that retrospective performance alone is insufficient to guarantee clinical impact, highlighting the importance of prospective and workflow-oriented validation studies to assess real-world utility and adoption (Aggarwal et al., 2021; Corbella; Srinivas; Cabitza, 2021).

Finally, the classification labels were derived solely from radiographic appearance, without individual-level confirmation by gold-standard vascular imaging techniques, such as Doppler ultrasonography or computed tomography angiography. While panoramic radiographs are known to reveal calcifications compatible with carotid atheromas, they do not allow direct assessment of plaque composition, degree of stenosis, or hemodynamic relevance (David et al., 2024; Saxena; Ng; Lim, 2019). Consequently, image-based labeling introduces an inherent degree of uncertainty and does not enable direct correlation between radiographic findings and clinical severity or patient outcomes. Incorporation of patient-matched vascular imaging would strengthen ground-truth definition and support more clinically nuanced performance evaluation.

In summary, the present results should be interpreted within the context of these limitations. While the proposed methodological framework demonstrates strong feasibility and diagnostic potential, additional evidence is required before clinical adoption. Larger and more diverse datasets, prospective validation studies, and integration with direct clinical reference standards are necessary before the approach can be considered for widespread clinical deployment or population-level screening programs.

6.6 Perspectives and Future Directions

Future research arising from the present work should prioritize the expansion and external validation of the proposed methodology to consolidate its clinical applicability and translational impact. A primary next step involves the construction of

larger multicenter datasets, incorporating panoramic radiographs acquired across diverse geographical regions, clinical environments, patient demographics, and imaging devices. Broader and more heterogeneous datasets enable more rigorous assessment of model generalization and are essential to mitigate domain bias and performance degradation when artificial intelligence systems are deployed in real-world clinical environments (Aggarwal et al., 2021; Zhou et al., 2021). Such expansion would further strengthen the robustness and reliability of the proposed approach under routine clinical conditions.

In parallel, prospective studies are essential to evaluate algorithm performance within real dental workflows. Unlike retrospective analyses, prospective study designs allow direct assessment of operational factors such as real-time image processing, clinician interaction with algorithmic outputs, referral compliance, delays between screening and confirmatory vascular assessment, and overall clinical acceptance. Prior work in medical artificial intelligence consistently emphasizes that prospective validation is a critical step for demonstrating true clinical impact beyond retrospective diagnostic accuracy (Aggarwal et al., 2021; Corbella; Srinivas; Cabitza, 2021).

Another important research direction concerns the integration of imaging-based predictions with electronic health records (EHRs). Combining radiographic findings with patient-specific clinical variables—such as age, sex, smoking status, diabetes, hypertension, lipid profile, and history of cardiovascular events—may enhance predictive performance and support the development of multimodal risk stratification models. Multimodal approaches have been widely recognized as a promising pathway toward personalized risk assessment and precision medicine, enabling more refined screening strategies and optimized referral thresholds (Johnson et al., 2021; Zhou et al., 2021).

The application of explainable artificial intelligence (XAI) represents another critical avenue for future development. Visualization techniques such as Gradient-weighted Class Activation Mapping (Grad-CAM) and related saliency-based methods can highlight image regions that most strongly influence model predictions, providing clinicians with intuitive explanations for automated alerts. Explainability has been identified as a key requirement for trust, validation, and safe adoption of AI systems in healthcare, supporting both clinical decision-making and systematic error analysis (Aggarwal et al., 2021; Zhou et al., 2021).

Finally, the ultimate translation of this research requires integration and evaluation within real clinical imaging pipelines. Embedding the trained model into dental imaging software or picture archiving and communication systems (PACS) would enable continuous background analysis of panoramic examinations, with automatic generation of alerts or annotations. Pilot deployment studies could then assess system responsiveness, computational performance on standard clinical hardware, impact on reporting workflows, and overall contribution to preventive cardiovascular care, as recommended in current guidelines for the clinical translation of AI-based diagnostic tools (Aggarwal et al., 2021; Corbella; Srinivas; Cabitza, 2021).

Collectively, these future directions aim to transition the present work from proof-of-concept algorithmic development toward clinically validated deployment, establishing dental panoramic imaging as a scalable and accessible platform for opportunistic vascular screening supported by artificial intelligence.

7 CONCLUSIONS

This dissertation presented the development and evaluation of a deep learning–based approach for the automatic detection of carotid atheromas on panoramic dental radiographs, using standardized regions of interest at the anatomical C3–C4 level and the MobileNetV2 architecture with a transfer learning strategy. The results demonstrated that the proposed method is capable of robustly and reliably identifying radiographic patterns associated with carotid calcifications. These findings confirm the feasibility of integrating artificial intelligence with dental radiology as a tool for opportunistic vascular risk screening.

The main objective of this work — to investigate the feasibility and performance of an automated system for carotid atheroma screening in routine panoramic radiographs — was fully achieved. The developed model showed high discriminative capability, maintaining an appropriate balance between detection of positive cases and exclusion of false-positive findings across different operational scenarios, both in balanced mode and in sensitivity-oriented screening settings. Statistical stability, demonstrated through bootstrap resampling analysis, reinforced the consistency of the results and enhanced confidence in the estimated performance metrics.

Among the principal scientific contributions of this dissertation are:

- The development of a standardized, anatomy-driven ROI pipeline (C3–C4) capable of reducing visual noise, increasing learning efficiency, and promoting model generalization;
- The successful application of a lightweight convolutional architecture (MobileNetV2), adapted to panoramic radiography through transfer learning and progressive fine-tuning, achieving competitive performance without reliance on complex or computationally intensive models;
- The integration of class-balancing strategies and realistic data augmentation, which enhanced sensitivity while maintaining control over specificity;
- A methodologically robust evaluation framework incorporating clinically relevant performance metrics, ROC and Precision–Recall analysis, systematic decision threshold selection, and quantification of statistical uncertainty via bootstrap resampling — an approach still infrequently adopted in this research field.

From a scientific and clinical perspective, this dissertation reinforces the role of panoramic dental radiography — a highly accessible and routinely acquired examination — as a complementary resource for opportunistic screening of vascular risk supported by artificial intelligence. The proposed methodology contributes to strengthening the interface between dentistry and preventive medicine, expanding the role of dental professionals in the early identification of systemic conditions associated with cerebrovascular disease. Its computational efficiency further supports practical implementation within conventional dental software platforms and imaging workstations, facilitating technology transfer to real-world clinical use.

In summary, this work demonstrates that the combination of anatomical standardization, lightweight convolutional neural networks, and rigorous methodological evaluation constitutes an effective strategy for automated detection of carotid atheromas on panoramic radiographs. By offering a technically feasible, statistically robust, and clinically relevant solution, this dissertation contributes to advancing artificial intelligence applications in preventive dentistry, with potential impact on population-level preventive strategies and future implementation in public health screening programs. While promising, these findings should be interpreted within the methodological scope of the present study and further validated through larger datasets and prospective clinical investigations.

8 REFERENCES

ABISSEGUE, Gisele *et al.* A systematic review of the epidemiology and the public health implications of stroke in Sub-Saharan Africa. **Journal of Stroke and Cerebrovascular Diseases**, v. 33, n. 8, p. 107733, 1 ago. 2024. https://doi.org/10.1016/j.jstrokecerebrovasdis.2024.107733

AGGARWAL, Ravi *et al.* Diagnostic accuracy of deep learning in medical imaging: a systematic review and meta-analysis. **npj Digital Medicine**, v. 4, n. 1, p. 65, 7 abr. 2021. https://doi.org/10.1038/s41746-021-00438-z

ALSAKAR, Yasmin M. *et al.* Multi-label dental disorder diagnosis based on MobileNetV2 and swin transformer using bagging ensemble classifier. **Scientific Reports**, v. 14, n. 1, p. 25193, 24 out. 2024. https://doi.org/10.1038/s41598-024-73297-9

ALZUBAIDI, Laith *et al.* Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. **Journal of Big Data**, v. 8, n. 1, p. 53, 2021. https://doi.org/10.1186/s40537-021-00444-8

AMITAY, Moshe *et al.* Deep convolution neural network for screening carotid calcification in dental panoramic radiographs. **PLOS Digital Health**, v. 2, n. 4, p. e0000081, 12 abr. 2023. https://doi.org/10.1371/journal.pdig.0000081

ANWAR, Syed Muhammad *et al.* Medical Image Analysis using Convolutional Neural Networks: A Review. **Journal of Medical Systems**, v. 42, n. 11, p. 226, 8 out. 2018. https://doi.org/10.1007/s10916-018-1088-1

ARZANI, Sarah *et al.* Detection of carotid artery calcifications using artificial intelligence in dental radiographs: a systematic review and meta-analysis. **BMC Medical Imaging**, v. 25, n. 1, p. 174, 19 maio 2025. https://doi.org/10.1186/s12880-025-01719-9

BUSHNELL, Cheryl *et al.* 2024 Guideline for the Primary Prevention of Stroke: A Guideline From the American Heart Association/American Stroke Association. **Stroke**, v. 55, n. 12, p. e344–e424, dez. 2024. https://doi.org/10.1161/STR.0000000000000482

CONSELHO NACIONAL DE SAÚDE. **Resolução nº 466**. Disponível em: <https://www.gov.br/conselho-nacional-de-saude/pt-br/atos-normativos/resolucoes/2012/resolucao-no-466.pdf/view>. Acesso em: 24 jan. 2026.

CONSELHO NACIONAL DE SAÚDE. **Resolução nº 510**. Disponível em: <https://www.gov.br/conselho-nacional-de-saude/pt-br/atos-normativos/resolucoes/2016/resolucao-no-510.pdf/view>. Acesso em: 24 jan. 2026.

CORBELLA, Stefano; SRINIVAS, Shanmukh; CABITZA, Federico. Applications of deep learning in dentistry. **Oral Surgery, Oral Medicine, Oral Pathology and Oral Radiology**, v. 132, n. 2, p. 225–238, 1 ago. 2021. https://doi.org/10.1016/j.oooo.2020.11.003

DAOLIO, Raul Muffato *et al.* Accuracy of duplex ultrasonography versus angiotomography for the diagnosis of extracranial internal carotid stenosis. **Revista do Colégio Brasileiro de Cirurgiões**, v. 51, p. e20243632, 2024. https://doi.org/10.1590/0100-6991e-20243632

DASHTI, Mahmood *et al.* Evaluation of accuracy of deep learning and conventional neural network algorithms in detection of dental implant type using intraoral radiographic images: A systematic review and meta-analysis. **The Journal of Prosthetic Dentistry**, v. 133, n. 1, p. 137–146, 1 jan. 2025. https://doi.org/10.1016/j.prosdent.2023.11.030

DAVID, Emanuele *et al.* Imaging of Carotid Stenosis: Where Are We Standing? Comparison of Multiparametric Ultrasound, CT Angiography, and MRI Angiography, with Recent Developments. **Diagnostics**, v. 14, n. 16, 5 ago. 2024. https://doi.org/10.3390/diagnostics14161708

DE ONOFRE, Niége Michelle Lazzari *et al.* Association between internal carotid artery calcifications detected as incidental findings and clinical characteristics associated with atherosclerosis: A dental volumetric tomography study. **European Journal of Radiology**, v. 145, p. 110045, 1 dez. 2021. https://doi.org/10.1016/j.ejrad.2021.110045

EFRON, B. AND TIBSHIRANI, R.J. **An Introduction to the Bootstrap.** Disponível em: <https://www.scirp.org/reference/referencespapers?referenceid=1879586>. Acesso em: 24 jan. 2026.

ER, Sezgin. **DENTEX CHALLENGE 2023**. Zenodo, , 9 abr. 2023. Disponível em: <https://zenodo.org/records/7812323>. Acesso em: 23 jan. 2026

ESTEVA, Andre *et al.* Deep learning-enabled medical computer vision. **npj Digital Medicine**, v. 4, n. 1, p. 5, 8 jan. 2021. https://doi.org/10.1038/s41746-020-00376-2

FELIPE, Beatriz Caio *et al.* Problem-based learning in dentistry: Diagnostic capability of dentists in the detection of calcified carotid artery atheroma on digital panoramic radiographs. **Research, Society and Development**, v. 9, n. 11, p. e70991110451–e70991110451, 1 dez. 2020. https://doi.org/10.33448/rsd-v9i11.10451

GARAGOLI, Fernando; MASSON, Walter; BARBAGELATA, Leandro. Association between elevated lipoprotein(a) levels and vulnerability of carotid atherosclerotic plaque: A systematic review. **Journal of Stroke and Cerebrovascular Diseases**, v. 33, n. 12, p. 108020, 1 dez. 2024. https://doi.org/10.1016/j.jstrokecerebrovasdis.2024.108020

GEIS, J. Raymond *et al.* Ethics of Artificial Intelligence in Radiology: Summary of the Joint European and North American Multisociety Statement. **Journal of the American College of Radiology: JACR**, v. 16, n. 11, p. 1516–1521, nov. 2019. https://doi.org/10.1016/j.jacr.2019.07.028

GOODFELLOW, Ian; BENGIO, Yoshua; COURVILLE, Aaron. **Deep Learning**. Cambridge, MA, USA: MIT Press, 2016.

HE, Kaiming *et al.* Deep Residual Learning for Image Recognition. *In*: 2016 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR). **2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)**. jun. 2016. Disponível em: <https://ieeexplore.ieee.org/document/7780459>. Acesso em: 24 jan. 2026 https://doi.org/10.1109/CVPR.2016.90

IKEDA, Naoki *et al.* Improvement of Detection Accuracy for Calcification Regions in Dental Panoramic Radiographs Using LVAT. *In*: 2024 INTERNATIONAL SYMPOSIUM ON INTELLIGENT SIGNAL PROCESSING AND COMMUNICATION SYSTEMS (ISPACS).

2024 International Symposium on Intelligent Signal Processing and Communication Systems (ISPACS). dez. 2024. Disponível em: <https://ieeexplore.ieee.org/abstract/document/10868471>. Acesso em: 23 jan. 2026 https://doi.org/10.1109/ISPACS62486.2024.10868471

JANISZEWSKA-OLSZOWSKA, Joanna *et al.* Carotid Artery Calcifications on Panoramic Radiographs. **International Journal of Environmental Research and Public Health**, v. 19, n. 21, p. 14056, 28 out. 2022. https://doi.org/10.3390/ijerph192114056

JIN, Weina *et al.* Guidelines and evaluation of clinical explainable AI in medical image analysis. **Medical Image Analysis**, v. 84, p. 102684, 1 fev. 2023. https://doi.org/10.1016/j.media.2022.102684

JOHNSON, Kevin B. *et al.* Precision Medicine, AI, and the Future of Personalized Health Care. **Clinical and Translational Science**, v. 14, n. 1, p. 86–93, 2021. https://doi.org/10.1111/cts.12884

KATI, Ezgi *et al.* Applying deep learning techniques to identify tonsilloliths in panoramic radiography. **Scientific Reports**, v. 15, n. 1, p. 24773, 9 jul. 2025. https://doi.org/10.1038/s41598-025-10489-x

KATS, Lazar *et al.* Atherosclerotic carotid plaque on panoramic radiographs: neural network detection. **International Journal of Computerized Dentistry**, v. 22, n. 2, p. 163–169, 2019.

KELLY, Christopher J. *et al.* Key challenges for delivering clinical impact with artificial intelligence. **BMC medicine**, v. 17, n. 1, p. 195, 29 out. 2019. https://doi.org/10.1186/s12916-019-1426-2

LECUN, Yann; BENGIO, Yoshua; HINTON, Geoffrey. Deep learning. **Nature**, v. 521, n. 7553, p. 436–444, maio 2015. https://doi.org/10.1038/nature14539

LEE, Sangyeon; KIM, Donghyun; JEONG, Ho-Gul. Detecting 17 fine-grained dental anomalies from panoramic dental radiography using artificial intelligence. **Scientific Reports**, v. 12, n. 1, p. 5172, 25 mar. 2022. https://doi.org/10.1038/s41598-022-09083-2

MANTA, Mihaela Daniela *et al.* The vertical topography of the carotid bifurcation – original study and review. **Surgical and Radiologic Anatomy**, v. 46, n. 8, p. 1253–1263, 1 ago. 2024. https://doi.org/10.1007/s00276-024-03404-y

MAYIL, Meltem; KESER, Gaye; PEKINER, Filiz. Clinical Image Quality Assessment in Panoramic Radiography. **Journal of Marmara University Institute of Health Sciences**, p. 1, 2014. https://doi.org/10.5455/musbed.20140610014118

MUREȘANU, Sorana; HEDEȘIU, Mihaela; IACOB, Liviu-Mihai. **Dataset for Automating Dental Condition Detection on Panoramic Radiographs**. Zenodo, , 22 maio 2025. Disponível em: <https://zenodo.org/records/15487430>. Acesso em: 23 jan. 2026

NELSON, John; VADDI, Anusha; TADINADA, Aditya. Can convolutional neural networks identify external carotid artery calcifications? **Oral Surgery, Oral Medicine, Oral Pathology and Oral Radiology**, v. 138, n. 1, p. 142–148, jul. 2024. https://doi.org/10.1016/j.oooo.2023.01.017

OMAROV, Murad *et al.* Automated Deep Learning-Based Detection of Early Atherosclerotic Plaques in Carotid Ultrasound Imaging. **medRxiv: The Preprint Server for Health Sciences**, p. 2024.10.17.24315675, 3 set. 2025.

PAKIZER, David *et al.* Diagnostics Accuracy of Magnetic Resonance Imaging in Detection of Atherosclerotic Plaque Characteristics in Carotid Arteries Compared to Histology: A Systematic Review. **Journal of Magnetic Resonance Imaging**, v. 61, n. 3, p. 1067–1093, 2025. https://doi.org/10.1002/jmri.29522

POSPELOV, B. A. *et al.* Morphological Changes in the Walls of Arteries of Different Basins at the Early Stages of Atherosclerosis Development. **Bulletin of Experimental Biology and Medicine**, v. 178, n. 5, p. 661–665, mar. 2025. https://doi.org/10.1007/s10517-025-06393-z

PRADOS-PRIVADO, María *et al.* Are Panoramic Images a Good Tool to Detect Calcified Carotid Atheroma? A Systematic Review. **Biology**, v. 11, n. 11, 20 nov. 2022. https://doi.org/10.3390/biology11111684

PRINS, Susanna R. *et al.* Life expectancy after an ischemic stroke or transient ischemic attack in older adults - the role of frailty. **Journal of Stroke and Cerebrovascular Diseases**, v. 34, n. 11, p. 108448, 1 nov. 2025. https://doi.org/10.1016/j.jstrokecerebrovasdis.2025.108448

ROC Radiologia. Disponível em: <https://rocradiologia.com.br/a-radiologia-odontologica-na-prevencao-do-avc/>. Acesso em: 5 jan. 2026.

RUSSAKOVSKY, Olga *et al.* **ImageNet Large Scale Visual Recognition Challenge**. arXiv, , 30 jan. 2015. Disponível em: <http://arxiv.org/abs/1409.0575>. Acesso em: 24 jan. 2026

SAATI, Samira; FOROOZANDEH, Maryam; ALAFCHI, Behnaz. Radiographic Characteristics of Soft Tissue Calcification on Digital Panoramic Images. **Pesquisa Brasileira em Odontopediatria e Clínica Integrada**, v. 20, p. e5053, 2020. https://doi.org/10.1590/pboci.2020.068

SAITO, Takaya; REHMSMEIER, Marc. The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets. **PLOS ONE**, v. 10, n. 3, p. e0118432, 4 mar. 2015. https://doi.org/10.1371/journal.pone.0118432

SANDLER, Mark *et al.* **MobileNetV2: Inverted Residuals and Linear Bottlenecks**. arXiv, , 21 mar. 2019. Disponível em: <http://arxiv.org/abs/1801.04381>. Acesso em: 24 jan. 2026

SANTOS, José Vinicius dos *et al.* Acidente Vascular Cerebral no Brasil: aspectos epidemiológicos da mortalidade no período de 2019 a 2023. **Brazilian Journal of Implantology and Health Sciences**, v. 7, n. 3, p. 1429–1439, 16 mar. 2025. https://doi.org/10.36557/2674-8169.2025v7n3p1429-1439

SAXENA, Ashish; NG, Eddie Yin Kwee; LIM, Soo Teik. Imaging modalities to diagnose carotid artery stenosis: progress and prospect. **BioMedical Engineering OnLine**, v. 18, n. 1, p. 66, 28 maio 2019. https://doi.org/10.1186/s12938-019-0685-7

SCHIROLI, Guido *et al.* Dental Panoramic Radiographs as Opportunistic Screening for Carotid Calcifications: A Case-Based Review. **Cureus**, 14 out. 2025. https://doi.org/10.7759/cureus.94585

SHUTTERSTOCK. **Cybermagician**. Disponível em: <https://www.shutterstock.com/pt/>. Acesso em: 5 mar. 2026.

SIMONYAN, Karen; ZISSERMAN, Andrew. **Very Deep Convolutional Networks for Large-Scale Image Recognition**. arXiv, , 10 abr. 2015. Disponível em: <http://arxiv.org/abs/1409.1556>. Acesso em: 24 jan. 2026

SONG, Yool Bin *et al.* Comparison of detection performance of soft tissue calcifications using artificial intelligence in panoramic radiography. **Scientific Reports**, v. 12, n. 1, p. 19115, 9 nov. 2022. https://doi.org/10.1038/s41598-022-22595-1

TZIOTZIOU, Aikaterini *et al.* Pressure- and flow-driven biomechanical factors associate with carotid atherosclerosis assessed by computed tomography angiography. **Atherosclerosis**, v. 408, p. 120415, 1 set. 2025. https://doi.org/10.1016/j.atherosclerosis.2025.120415

VINAYAHALINGAM, Shankeeth *et al.* Detection of carotid plaques on panoramic radiographs using deep learning. **Journal of Dentistry**, v. 151, p. 105432, 1 dez. 2024. https://doi.org/10.1016/j.jdent.2024.105432

WEI, Xufang *et al.* A systematic review and meta-analysis of clinical efficacy of early and late rehabilitation interventions for ischemic stroke. **BMC Neurology**, v. 24, n. 1, p. 91, 8 mar. 2024a. https://doi.org/10.1186/s12883-024-03565-8

WEI, Yao *et al.* Real-time carotid plaque recognition from dynamic ultrasound videos based on artificial neural network. **Ultraschall in der Medizin - European Journal of Ultrasound**, v. 45, n. 5, p. 493–500, out. 2024b. https://doi.org/10.1055/a-2180-8405

WORLD HEALTH ORGANIZATION. **Global Atlas on Cardiovascular Disease Prevention and Control**. Geneva, Switzerland: *[S.n.]*.

WORLD HEALTH ORGANIZATION. **Leading causes of death**. Disponível em: <https://www.who.int/data/gho/data/themes/mortality-and-global-health-estimates/ghe-leading-causes-of-death>. Acesso em: 23 jan. 2026.

WORLD HEALTH ORGANIZATION. **The top 10 causes of death**. Disponível em: <https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death>. Acesso em: 23 jan. 2026.

WORLD MEDICAL ASSOCIATION. World Medical Association Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Participants. **JAMA**, v. 333, n. 1, p. 71–74, 7 jan. 2025. https://doi.org/10.1001/jama.2024.21972

YOSINSKI, Jason *et al.* **How transferable are features in deep neural networks?** arXiv, , 6 nov. 2014. Disponível em: <http://arxiv.org/abs/1411.1792>. Acesso em: 24 jan. 2026

YOUDEN, W. J. Index for rating diagnostic tests. **Cancer**, v. 3, n. 1, p. 32–35, jan. 1950. https://doi.org/10.1002/1097-0142(1950)3:1<32::AID-CNCR2820030106>3.0.CO;2-3

ZHOU, S. Kevin *et al.* A Review of Deep Learning in Medical Imaging: Imaging Traits, Technology Trends, Case Studies With Progress Highlights, and Future Promises. **Proceedings of the IEEE**, v. 109, n. 5, p. 820–838, maio 2021. https://doi.org/10.1109/JPROC.2021.3054390

ZHU, Ying *et al.* The application of the nnU-Net-based automatic segmentation model in assisting carotid artery stenosis and carotid atherosclerotic plaque evaluation. **Frontiers in Physiology**, v. 13, 6 dez. 2022. https://doi.org/10.3389/fphys.2022.1057800