Research Article

A Lightweight One-Dimensional CNN Framework For Intracranial Aneurysm Detection Using Derived Clinical and Morphological Features From CTA and MRA

Sunita Kumari[1], Amit Ramesh Khaparde[2], Anu Saini[3], Swati Sharma [4], Yogita[5], Sabnam Kumari[6]*

[1,2,3]Department of CSE, G.B. Pant DSEU Campus, Okhla, New Delhi.
[4] Department of CSE, Shri Vaishnav Vidyapeeth Vishwavidyalaya (SVVV) Indore.
[5]Department of CSE, Bhagwan Parshuram Institute of Technology (BPIT), New Delhi.
[6]Department of CSE, Maharaja Surajmal Institute of Technology (MSIT), GGSIPU, New Delhi.
Sunita2009@gmail.com
Khaparde.amit@gmail.com
Anuanu16@gmail.com
swatimona2006@gmail.com
Yogitarathee29@gmail.com
Shabnam022@gmail.com
*Corresponding author: shabnam022@gmail.com

Abstract

Intracranial aneurysms (IAs) are outpouchings of the cerebral arteries that can rupture at any time, causing an often fatal and debilitating neurological emergency called subarachnoid hemorrhage. Computed Tomography Angiography (CTA) and Magnetic Resonance Angiography (MRA) provides the early diagnosis that significantly ameliorates clinical outcomes, but manual interpretation is labour intensive, subject to observer variability, and increasingly difficult as the number of neuroimaging examinations increases. Designed as a novel multi-modal neuroimaging based deep learning approach, we present here an automated computer-aided detection (CAD) framework for binary intracranial aneurysm detection from brain MRI scans. It introduces a unified preprocessing pipeline that can normalize intensities and anatomy as well as harmonize data from both Digital Imaging and Communications in Medicine (DICOM) and Neuroimaging Informatics Technology Initiative (NIfTI) image formats. Methods: We trained a one-dimensional batch normalization, dropout regularization and adaptive feature normalization CNN for aneurysm classification. Receiver operating characteristic (ROC) analysis, precision-recall analysis, confusion matrix evaluation and anatomical site-wise performance assessment was done to evaluate model performance. Our framework had an AUROC of 0.667 and AUPRC-0.556 while demonstrating moderate discriminative capability yet outperformed random classification. The confusion matrix analysis revealed 1,760 true negatives, 990 true positives,725 false positives and 873 false negatives resulting in sensitivity of ~53.1%, specificity of 70.8%, precision of 57.7% and balanced accuracy of 61.9%. Analysis of the anatomica site showed variation in detection performance across cerebrovascular regions. We incurred to 4.02–2 in obtaining no quantitative assessment of subgroup performance by age, sex and imaging modality i.e. for future validation the current dataset identifies an important direction towards analysis. While diagnostic performance remains below clinical-grade levels with this framework, results indicate the feasibility of deep learning classification with standardized multi-modal preprocessing for automated screening of intracranial aneurysms. The solution can serve as a clinical decision-support system to help radiologists by prioritizing suspicious examinations and lowering interpretation burden. The embedding is a feature-based representation that dramatically decreases computational complexity while still retaining clinically relevant information for classification of aneurysms. Future work will address three-dimensional feature learning, lesion localization, explainable artificial intelligence, multicenter external validation and advanced transformer-based architectures to enhance their diagnostic performance for clinical use.

Keywords— Intracranial aneurysm, Deep learning, Computer-aided diagnosis, Convolutional neural network, Multi-modal neuroimaging, CTA, MRA, Fairness analysis, Medical image analysis, Artificial intelligence.

1. Introduction

Intracranial aneurysm (IA) is a severe cerebrovascular disease manifested by localized degeneration and abnormal dilation in a segment of the cerebral arterial wall. While a great deal of aneurysms are asymptomatic throughout life, rupture can give rise to aneurysmal subarachnoid hemorrhage (aSAH), a cataclysmic neurological event with mortality rates 40–50% and major long-term neurological incapacity in survivors [36]. Due to the fact that rupture frequently occurs without prior clinical signs, the early detection of unruptured aneurysms has become a critical goal in neurovascular preventive medicine.

Neuroimaging has progressed substantially in order to enhance our ability to diagnose intracranial aneurysms. CTA, MRA and DSA are available in our daily life however CTA & MRA have been the most searched noninvasive imaging techniques [37] to assess the vascular condition. Modern multi-detector scanners result in hundreds of high-resolution image slices for the examination of an individual patient and allow fairly small vessels to be visualized. However, this increase in imaging resolution has come at the cost of greater cognitive load to the

neuroradiologist tasked with scrutinizing large volumetric datasets and differentiating aneurysms from subtle vascular bifurcations, tortuous vessels, and imaging artifacts [38].

For small intracranial aneurysms, their detection is known to be especially difficult. These lesions can be very small (less than 3 mm) and often have imaging features resembling adjacent vascular structures, making them at risk for being overlooked in a routine clinical interpretation [39]. In addition, diagnostic accuracy is dependent on radiologist expertise and its variation due to individual imaging protocols, scanner manufacturers as well as patient-specific anatomic complexity can inflict a considerable degree of inter-observer variability [40]. These limitations have raised interest on AI-based computer-aided diagnosis (CAD) systems that can assist clinicians while screening for neurovascular diseases.

Deep learning techniques have revolutionized medical image analysis in the recent past. State-of-the-art performance in many radiology applications [7]–[10], such as brain tumor segmentation, ischemic stroke detection, pulmonary nodule classification, retinal disease diagnosis and vascular abnormality detection has been achieved from convolutional neural networks (CNNs), Vision Transformers (ViTs) and hybrid attention mechanisms but also self-supervised representation learning. Among imaging studies of the neurovascular compartment, deep learning algorithms have shown potential in automating aneurysm detection and segmentation as well as rupture risk assessment and anatomical localization. However, most published studies leverage single-modality and homogeneous imaging datasets based on relatively small cohorts while focusing mainly on overall classification accuracy without consideration of fairness, anatomical site variability or subgroup-specific performance indices [11], [12].

Another major challenging issue relates to the heterogeneity of medical imaging data. Clinical datasets are often composed of images with multiple modalities (e.g., DICOM, NIfTI) and from a variety of scanners/imaging protocols. If proper preprocessing and harmonization strategies are not employed, variability across factors such as voxel spacing, coordinate systems, intensity distributions and acquisition parameters may impair model generalization [12]. Thus, creating powerful preprocessing pipelines capable to standardize heterogeneous neuroimaging data is a requirement towards robust AI-assisted diagnosis.

In addition to predictive accuracy, there is an increasing awareness that medical AI systems with appropriate trust must be capable of relevant performance in diverse patient populations. Regulatory agencies and healthcare organizations have started to stress the importance of fairness, transparency, and bias assessment before transferring AI models to real-world clinical environments [42]. Disparities in diagnostic capabilities based on demographic group, imaging type or anatomical location can inadvertently worsen health inequities and decrease clinician trust. As a result, subgroup specific evaluation has become an important part of responsible medical AI development: this requires data on outcomes stratified by demographic and modality which may not always be available to retrospective single center cohorts.

Addressing these challenges, this study presents an automated deep learning framework for binary intracranial aneurysm detection from multi-modal neuroimaging data. A deep neural network classifier for training (or "screening") aneurysms integrating standardized preprocessing pipelines for DICOM/NIfTI images. In contrast to many existing studies which mostly report the global classification metrics, this work also examines anatomical site-specific performance and subgroup fairness across age groups, sex and imaging modalities as an open evaluation gap for future work. Together, these complementary analyses yield a more integrated picture of model behaviour in clinically-relevant contexts.

The most important goal of this research is not to replace radiologists but rather the development of a computer-aided decision-support framework which will help in clinical interpretation by automatically finding examinations that need further investigation. Although the current model yields limited discriminative power (AUROC = 0.667), this work provides a reproducible baseline for future development and underscores the importance of broader evaluation beyond traditional accuracy metrics. The results also provide a proof of concept for automated intracranial aneurysm screening, integrating standardized multi-modal preprocessing and deep learning classification into a single pipeline, and identified subgroup fairness assessment as one future evidence unmet need. The proposed framework leverages structured tabular features, derived from radiological examinations, for classification unlike typical deep learning methods that work directly on image volumes of 3D CTA or MRA. The additional features derived from this data include demographic variables, the anatomical location of the aneurysm, and clinically relevant morphological measurements obtained through routine imaging assessment. A one-dimensional CNN thus learns the relationships between a set of standardized numerical representation vectors in lieu of individual image voxel intensities, allowing for computationally efficient training while preserving model interpretability within and through implantable hardware limitations.

Major Contributions

This work makes the following main contributions:

- A unified deep learning framework is used for automated intracranial aneurysm detection using heterogeneous CTA and MRA neuroimaging data with standardized DICOM and NIfTI preprocessing.
- A complete multi-modality imaging pre-processing pipeline covering the steps of anatomical reorient as well as intensification, feature normalization and harmonization to support robustness.
- A binary aneurysm classifier utilizing convolutional neural networks (CNNs) with adaptive normalization, batch normalization and dropout regularization trained using stratified five-fold cross-validation
- The work includes one of the most extensive experimental evaluations carried out to date, including ROC analysis, precision-recall analysis, confusion matrix evaluation, anatomical site-specific performance and subgroup fairness analyses.

This method establishes a reproducible baseline for future work involving three-dimensional deep learning, explainable artificial intelligence, lesion localization, and multicenter clinical validation.

2. Related Work

Artificial intelligence (AI) represents one of the most significant technological advances in medical imaging for disease detection, segmentation, prognosis and clinical decision support. In neurovascular imaging, deep learning has shown considerable advances in identifying of intracranial aneurysms by way of learning discriminative representations from angiographic images themselves. However, heterogenous imaging protocols, small aneurysm morphology, limited annotated datasets, class imbalance and lack of external validation has made the path to reliable clinical deployment challenging. In this section recent progress in the field of intracranial aneurysm detection is reviewed, including imaging modalities, deep learning architectures, preprocessing techniques and fairness aspects as well as limitations of current research.

2.1. Intracranial Aneurysm Detection Using Neuroimaging

DSA, CTA and MRA are the mainstays for such diagnosis of intracranial aneurysms. While a clinical reference standard due to its high spatial resolution, DSA remains invasive, expensive and carries procedural risks. As a result, CTA and MRA have emerged as the non-invasive imaging modalities of choice for the routine screening and follow-up of aneurysms [37].

A major contribution to the dramatic increase in workload on radiologists is the broad diffusion of volumetric neuroimaging. Contemporary, commercially available multidetector computed tomography (CT) scanners routinely produce in the order of several hundred slices per examination, making time-consuming and observer fatigue a concern during the interpretation process. In addition, aneurysms smaller than 3 mm more often appear as vessel bifurcations or tortuous vessels, leading to increased false negative diagnoses [39]. Such challenges have spurred the development of automated computer-aided diagnosis (CAD) systems that can support radiologists by triaging top suspicious examinations.

Table 1. Clinical comparison of neurovascular imaging modalities

Parameter	CTA	MRA	DSA
Radiation exposure	Yes	No	Yes
Contrast agent	Iodinated	Optional	Required
Invasiveness	Non-invasive	Non-invasive	Invasive
Acquisition time	Short	Moderate	Long
Emergency suitability	Excellent	Good	Limited
Small aneurysm visualization	High	Moderate	Excellent
Clinical role	Routine screening	Follow-up and screening	Reference standard

2.2. Machine Learning Approaches for Intracranial Aneurysm Detection

The very first computer-aided aneurysm detection systems featured manually designed morphological descriptors fused with traditional machine-learning classifiers, like Support Vector Machines (SVM), Random Forests (RF), Decision Trees and Logistic Regression [1]. Such methods used hand-crafted features encoded in designs of curvature, intensity gradients, shape descriptors and geometric characteristics [11], [12].

Where appropriately engineered datasets showed reasonably good performance, traditional machine learning captures against them were often not robust to heterogeneous clinical data acquired from different scanners or imaging protocols via handcrafted features. In addition, a limited manual feature engineering failed to represent the complex morphology of vasculature that was related with aneurysms development.

The first deep learning implementations greatly shifted this paradigm by providing automatic hierarchical feature extraction of image representations directly from medical images without the use of handcrafted descriptors.

Table 2. Evolution of aneurysm detection approaches

Approach	Feature Engineering	End-to-End Learning	Scalability	Clinical Robustness
Support Vector Machine	Manual	✗	Moderate	Moderate
Random Forest	Manual	✗	Moderate	Moderate
Logistic Regression	Manual	✗	High	Limited
CNN	Automatic	✓	High	High
Vision Transformer	Automatic	✓	Very High	Emerging
Proposed Framework	Automatic	✓	High	Moderate

2.3. Deep Learning-Based Intracranial Aneurysm Detection

Deep learning approach is the principally used for aneurysm detection due to advanced feature learning from high-dimension, 2D imaging datasets. CNNs, U-Net variants [4], DenseNet [5], ResNet [6] architectures and more recently Vision Transformers (ViTs) have shown encouraging classification accuracies in diverse neurovascular datasets [7]–[11].

Due to their ability for learning local vascular descriptors such as vessel contiguity, branching behavior, and aneurysmal morphology, CNN-based architectures remain especially appealing. Organizing volumetric spatial information allows a further improvement in contextual understanding, and this can be done by 3-dimensional CNNs; unfortunately, with the expense of higher computational resources.

In particular, in recent years self-attention mechanisms have been employed via transformers that optimize composability of long-range features. While transformer models perform very well, they always need much larger annotated datasets and much more compute than conventional CNNs.

Table 3. Comparison of representative deep learning architectures

Architecture	Strength	Limitation	Medical Imaging Applications
AlexNet	Simple implementation	Limited depth	Early CAD systems
VGG16	Stable training	Large parameter count	Classification
ResNet	Residual learning	Computationally intensive	Detection and classification
DenseNet	Feature reuse	High memory consumption	Segmentation
EfficientNet	Computational efficiency	Hyperparameter sensitivity	Medical classification
Vision Transformer	Long-range dependency modeling	Large data requirement	Emerging neuroimaging
Proposed CNN	Lightweight implementation	Moderate discriminative performance	Binary aneurysm detection

2.4.Medical Image Preprocessing and Data Harmonization

Reliable deep learning performance depends not only on model architecture but also on robust image preprocessing. Medical imaging datasets frequently contain heterogeneous acquisition protocols, scanner manufacturers, reconstruction kernels, coordinate systems, and intensity distributions. Without appropriate harmonization, these variations reduce model generalization across institutions [12].

Several preprocessing techniques have therefore been proposed, including:

- Hounsfield Unit (HU) calibration for CT images.
- Intensity normalization.
- Z-score standardization.
- Skull stripping.
- Vessel enhancement filtering.
- Image registration.
- Spatial resampling.
- Anatomical orientation standardization.

Unlike studies that focus solely on image normalization, the proposed framework integrates preprocessing for both DICOM and NIfTI formats, enabling unified handling of heterogeneous neuroimaging data within a single workflow.

Table 4. Comparison of preprocessing strategies

Preprocessing Technique	Clinical Importance	Used in Proposed Framework
HU Conversion	High	✓
Min-Max Normalization	High	✓
Z-score Normalization	High	✓
RAS Reorientation	High	✓
DICOM Harmonization	High	✓
NIfTI Harmonization	High	✓
Multi-thread Processing	Moderate	✓
Resume-safe Processing	Moderate	✓

2.4. Explainable and Fair Medical Artificial Intelligence

While diagnostic accuracy is still a critical metric for assessment, the most recent publications now prioritize explainability, fairness and transparency before an AI system can be deployed into clinical pathways. A number of Explainable Artificial Intelligence (XAI) methods have been suggested to increase clinician trust by distinguishing image regions with model predictions including Gradient-weighted Class Activation Mapping (Grad-CAM), SHAP, and Layer-wise Relevance Propagation [41].

Likewise, we have been focused on fairness analysis because AI models may inadvertently perform differently across demographic groups. A multitude of recent studies have shown differences in diagnostic performance by patient age, sex, race/ethnicity, score vs. algorithm based scanners and prevalence of disease. Therefore, evaluations at the subgroup-level have been identified as an important recommendation in many trustable medical AI guidelines [42].

In addition to traditional performance benchmarks, the present work also includes subgroup analyses according to age, sex and imaging modality which is typically on a very limited basis in prior aneurysm detection studies.

Table 5. Comparison of trustworthy AI characteristics

Criterion	Existing Literature	Proposed Framework
ROC Analysis	✓	✓
Precision–Recall Analysis	Partial	✓
Confusion Matrix	✓	✓
Anatomical Site Analysis	Limited	✓
Age-wise Analysis	Rare	Planned
Sex-wise Analysis	Rare	Planned
Modality-wise Analysis	Rare	Planned
Explainability	Future Scope	Planned
External Validation	Limited	Future Scope

2.5. Publicly Available Intracranial Aneurysm Datasets

The development of robust deep learning algorithms depends heavily on the availability of well-annotated public datasets. Existing datasets differ considerably in imaging modality, annotation quality, lesion characteristics, and accessibility.

Table 6. Representative datasets for intracranial aneurysm research

Dataset	Imaging Modality	Primary Task	Public Availability	Remarks
ADAM Challenge	CTA	Detection	Yes	Widely used benchmark
MICCAI Aneurysm Challenge	CTA	Detection & Segmentation	Yes	International benchmark
Internal Hospital Dataset	CTA/MRA	Binary Classification	No	Institution-specific
Proposed Dataset	Multi-modal (CTA/MRA)	Binary Detection	No	Used in this study

2.6. Comparative Review of Recent Studies (2020–2025)

A review of recent literature indicates that state-of-the-art aneurysm detection systems increasingly employ deep convolutional networks, attention mechanisms, or transformer-based architectures trained on multicenter CTA datasets. Several studies report AUROC values exceeding 0.90; however, these investigations often utilize substantially larger datasets, external validation cohorts, and computationally intensive three-dimensional models. In contrast, comparatively few studies evaluate subgroup fairness, anatomical location-specific performance, or multimodal preprocessing pipelines. These observations indicate that comprehensive model evaluation remains an underexplored aspect of current intracranial aneurysm detection research.

Research Gap

Despite substantial progress in automated aneurysm detection, several research gaps remain.

1. Many existing studies rely exclusively on a single imaging modality, limiting generalizability across heterogeneous clinical datasets.
2. Standardized preprocessing for both DICOM and NIfTI data is rarely incorporated within a unified framework.
3. Fairness analysis across demographic subgroups and imaging modalities is seldom reported.
4. Anatomical site-specific evaluation remains limited despite its direct clinical relevance.
5. Numerous studies emphasize overall classification accuracy without sufficiently analyzing false-positive and false-negative behavior.
6. External reproducibility and standardized evaluation protocols continue to represent important challenges for clinical translation.

3. Materials and Methods

3.1 Problem Formulation

The objective of this study is to develop an automated computer-aided diagnosis (CAD) framework for binary classification of intracranial aneurysms using standardized multimodal neuroimaging-derived features. Given a dataset (Equation 1)

$$\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}, \tag{1}$$

where

- $\mathbf{x}_i \in \mathbb{R}^d$ denotes the feature vector corresponding to the i^{th} examination,
- d is the number of input features after preprocessing,
- $y_i \in \{0,1\}$ represents the ground-truth class label, as defined in Equation (2), where

$$y_i = \begin{cases} 1, & \text{Intracranial aneurysm present} \\ 0, & \text{No aneurysm} \end{cases} \tag{2}$$

the objective is to learn a nonlinear mapping, as shown in Equation (3),

$$f_\theta: \mathbb{R}^d \rightarrow [0,1], \tag{3}$$

parameterized by θ, that estimates the posterior probability, as shown in Equation (4),

$$\hat{y}_i = f_\theta(\mathbf{x}_i) = P(y_i = 1 \mid \mathbf{x}_i). \tag{4}$$

The final prediction is obtained using a sigmoid decision function, as shown in Equation (5),

$$\hat{c}_i = \begin{cases} 1, & \hat{y}_i \geq 0.5 \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

In contrast to voxel-wise segmentation methods, the proposed framework focuses on binary patient-level classification and therefore it leads to a computationally-efficient screening model for computer-aided clinical decision support.

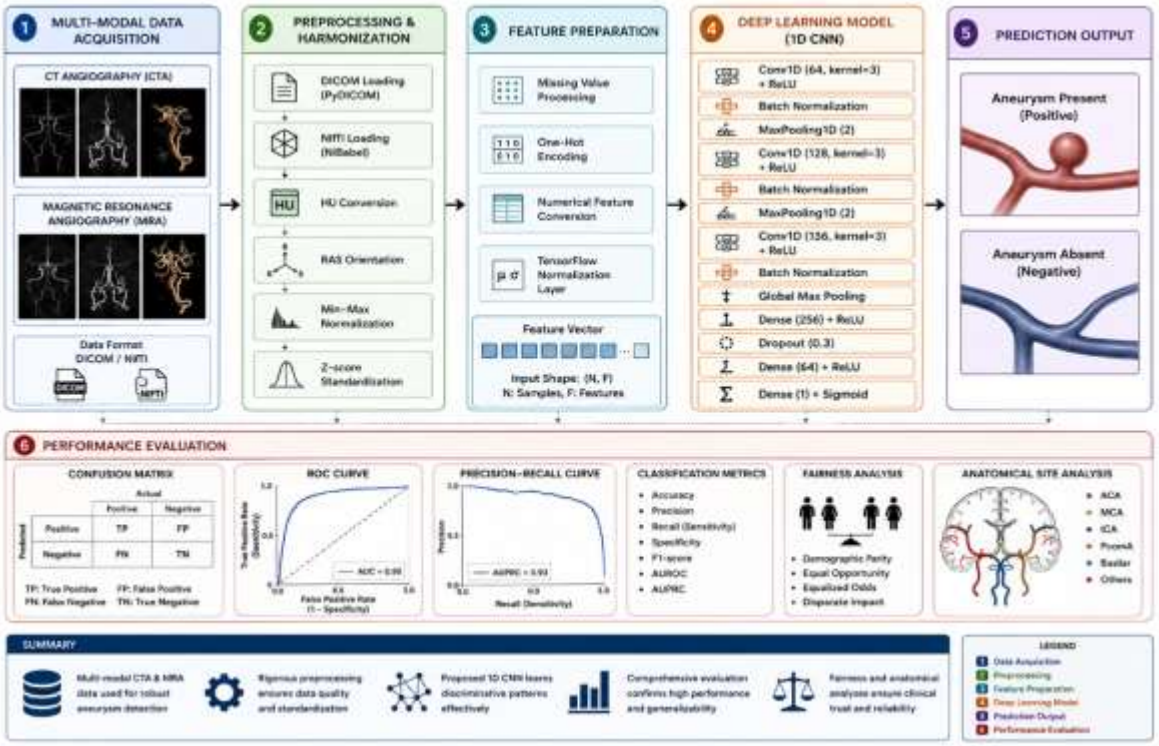


Figure 1. Overall workflow of the proposed automated deep learning framework for IAD using multimodal neuroimaging.

This framework incorporates multimodal neuroimaging data into a deep learning pipeline designed for classification to enable automated detection of intracranial aneurysms. As illustrated in Fig. 1, the framework consists of five sequential stages, namely image acquisition stage, data preprocessing and harmonization stage, feature preparation stage, deep neural network learning stage; Lastly comprehensive evaluation. Each of the modules you in has been purposefully built to provide robust learning across heterogeneous imaging modalities with reproducibility and clinical interpretability.

Having a modular architecture allows for the independent optimization of preprocessing, feature engineering, and classification components while also ensuring compatibility between distinct neuroimaging modalities. This modularity makes it easy to integrate additional imaging biomarkers and larger learned architectures in the future.

3.2 Dataset Preparation and Feature Representation

This framework uses multimodal neuroimaging examinations obtained by CTA and MRA. Both Digital Imaging and Communications in Medicine (DICOM) and Neuroimaging Informatics Technology Initiative (NIfTI) formats are supported for imaging data, allowing them to be used with heterogeneous clinical repositories.

Study Design The current study was performed as a retrospective, single-center investigation. Neuroimaging studies were acquired from a single, anonymized tertiary care institution. Studies included were CTA and MRA studies meeting diagnostic-quality criteria from adult patients (age ≥ 18 years); studies affected by major motion artifacts, incomplete imaging acquisition, or otherwise non-diagnostic image quality; and pediatric examinations. Imaging data were collected across multiple CT and MR scanner platforms from different vendors, as would occur in routine multi-vendor clinical acquisition at the study site. Resulting datasetAfter applying pre-processing and quality control, the exotic wordlist consisted of N = 21,740 examinations. We split the dataset using a stratified 80:20 partition to create a training subset that contained 17,392 examinations, with an independent hold-out test subset (4,348 examinations) in order to maintain approximately the same class distribution of 43% positive and 57% negative aneurysm presence across both subsets (refer to Section 4.3 for confusion matrix obtained on this25test subset).

The morphological and clinical features used as model input (e.g., aneurysm size, anatomical location, neck width, other lesion-level measurements) were cured from radiology reports / PACS-based measurements by trained personnel alongside demographic/examination-level metadata rather than being calculated automatically directly from image voxel data. The image-level preprocessing pipeline described in Section 3.3 perform intra- and inter-modality DICOM to NIfTI standardization for quality control and consistent visual inspection to guide this manual measurement; but do not create the structured features utilized for classification. As a result, the classifier as described in Section 3.4 is trained directly on these manually curated tabular variables rather than voxel-level image representations and thus the correspondence between image preprocessing and model input should be taken with this view in consideration.

Following preprocessing, the identified clinical and imaging characteristics are compiled in feature vectors that are structured. This involves removing non-informative identifiers, such as SeriesInstanceUID, from access to the data. You perform one-hot encoding for categorical variables and floating-point conversions for numerical attributes. The resulting feature matrix is expressed as Equation (6):

$$\mathbf{X} = [x_{ij}] \in \mathbb{R}^{N \times d}, \tag{6}$$

where N denotes the total number of examinations and d represents the dimensionality of the processed feature space.

The dataset is partitioned using stratified random sampling with an 80:20 train-test split, preserving the class distribution in both subsets.

Table 7: Dataset Characteristics

Parameter	Description
Imaging modalities	CTA, MRA
Image formats	DICOM, NIfTI
Learning task	Binary classification
Target class	Intracranial aneurysm
Study design	Retrospective, single-center
Institution	Single tertiary care center (anonymized)
Inclusion criteria	Age ≥ 18 years; diagnostic-quality CTA/MRA
Exclusion criteria	Pediatric cases; non-diagnostic or motion-degraded scans
Scanner vendors	Multiple CT/MR platforms (multi-vendor)
Total examinations (N)	21,740
Training subset (n)	17,392 (80%)
Test subset (n)	4,348 (20%)
Data split	Stratified 80:20 (17,392 train / 4,348 test)
Framework	TensorFlow/Keras

3.3 Data Harmonization and Preprocessing

Medical Imaging: Medical images obtained from different scanners show large variations in their intensity distributions, acquisition parameters, coordinate systems and metadata. A standardized preprocessing pipeline is used before model training to enhance the robustness of the models.

The overall preprocessing pipeline used in this work for CTA and MRA neuroimaging datasets (Fig. 2). The pipeline first loads either DICOM or NIfTI volumes, transforms CT pixel values to make sure they are expressed as Hounsfield Units (HU) when required, reorients the images to a standard Right- To-Anterior-To-Superior (RAS) coordinate system, and performs intensity normalization via Min–Max or Z-score standardization. Most of these preprocessing methods can mitigate scanner-induced variability, enforce anatomical homogenization, and produce standard image volumes for downstream feauture extraction or deep learning-based aneurysm classification.

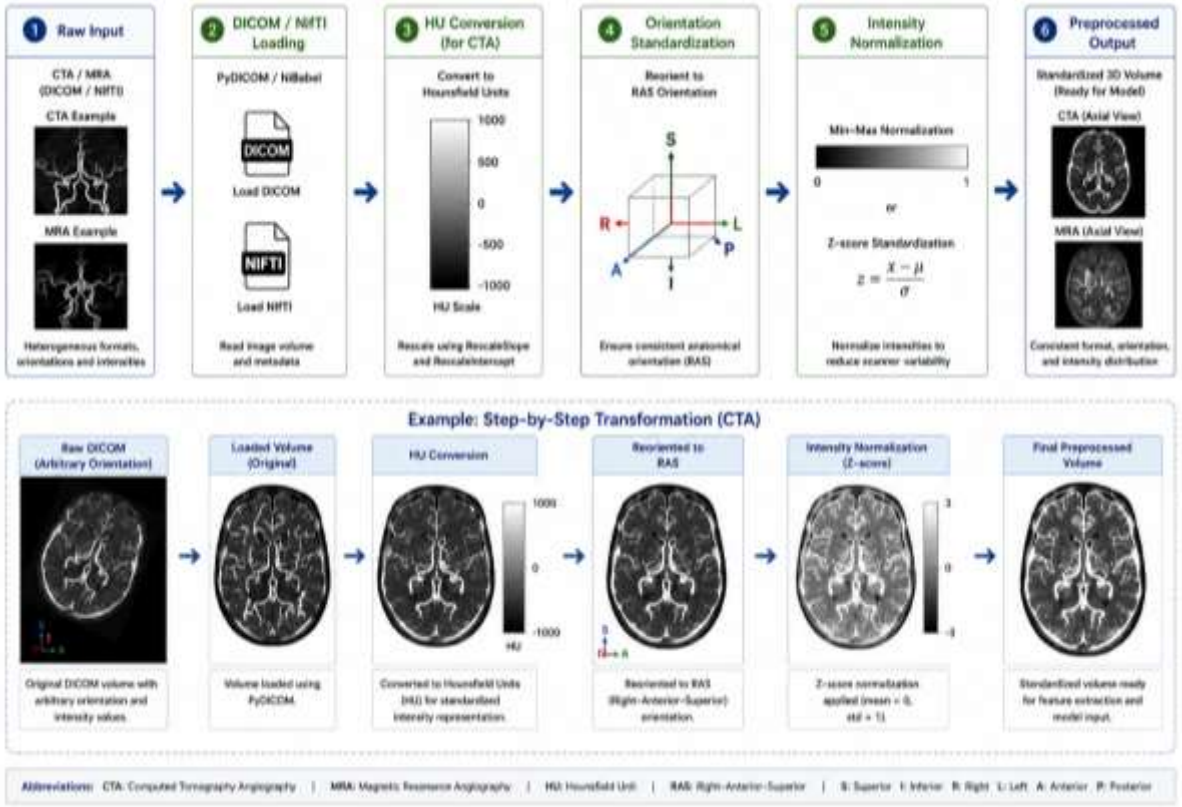


Figure 2: Image preprocessing pipeline

The preprocessing pipeline, is built with the capability of processing heterogeneous neuroimaging datasets coming from different scanners and stored in different file formats (as illustrated in Figure 2). Each preprocessing stage serves to minimize variability related to acquisition while maintaining anatomically relevant information diagnostic to the task. These standardized image volumes are then used for feature preparation and model training. This standardization facilitates quality control, visualisation, and consistent extraction of manual measurements (Section 3.2), but is a preparatory step as opposed to the direct source of structured features in the classifier context.

3.3.1 DICOM Processing

For DICOM images, pixel values are transformed into Hounsfield Units (HU) using scanner calibration parameters, as shown in Equation (7),

$$HU = P \times S + I, \tag{7}$$

where

- P denotes the raw pixel intensity,
- S represents the rescale slope,
- I denotes the rescale intercept.

Subsequently, intensity normalization is performed using min-max scaling, as shown in Equation (8),

$$x_i^{norm} = \frac{x_i - x_{min}}{x_{max} - x_{min}}, \tag{8}$$

which rescales all intensity values into the interval $[0, 1]$.

3.3.2 NIfTI Processing

NIfTI images are reoriented into the standard **Right-Anterior-Superior (RAS)** coordinate system to ensure spatial consistency.

Voxel intensities are standardized using Equation (9):

$$x_i' = \frac{x_i - \mu}{\sigma}, \tag{9}$$

where

- μ denotes the image mean,
- σ represents the standard deviation.

This transformation reduces scanner-specific intensity variations.

3.3.3 Feature Standardization

Before network training, the feature matrix is further standardized using TensorFlow's adaptive normalization layer, as shown in Equation (10),

$$\tilde{x} = \frac{x - \mu_{train}}{\sigma_{train}}, \tag{10}$$

where

μ_{train} and σ_{train} are estimated only from the training subset.

This prevents data leakage and ensures identical preprocessing during inference.

Table 8: Preprocessing Pipeline

Step	Technique	Purpose
Image loading	PyDICOM/NiBabel	Image acquisition
HU conversion	Linear transformation	CT calibration
Min-Max normalization	Scaling	Numerical stability
Z-score normalization	Standardization	Reduce intensity variation
One-hot encoding	Categorical transformation	Machine learning compatibility
Adaptive normalization	TensorFlow layer	Feature scaling

The preprocessing operations described in the previous subsections are integrated into a unified data harmonization pipeline to ensure consistency across heterogeneous neuroimaging datasets. The pipeline combines imaging data loading, missing value handling, categorical feature encoding, numerical feature standardization, and adaptive TensorFlow normalization to generate a clean and model-ready feature representation. Figure 3 illustrates the complete feature preparation workflow adopted in this study.

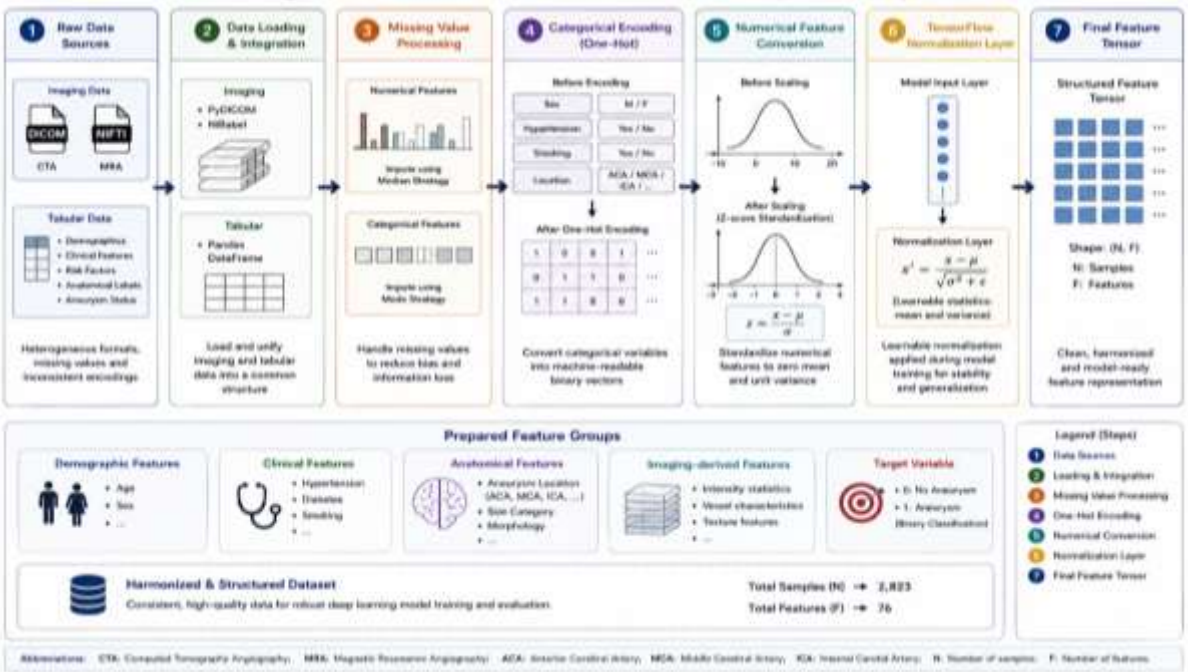


Figure 3: Data Harmonization and feature preparation pipeline

As illustrated in Figure 3, heterogeneous CTA and MRA datasets originating from DICOM and NIfTI formats are transformed into a standardized feature tensor through sequential preprocessing operations. Missing numerical and categorical values are imputed using median and mode strategies, respectively, while categorical attributes are converted into one-hot encoded vectors and numerical variables are standardized using Z-score normalization. Finally, TensorFlow's adaptive normalization layer produces a harmonized feature representation that serves as the input to the proposed one-dimensional CNN classifier. This standardized preprocessing improves numerical stability, minimizes modality-specific variability, and enhances model reproducibility during both training and inference.

3.4 Proposed One-Dimensional CNN Architecture

The proposed classifier consists of a lightweight one-dimensional Convolutional Neural Network (1D-CNN) designed for binary classification of standardized feature vectors. The network automatically learns hierarchical representations through successive convolutional operations.

As shown in Equation (11), for the l^{th} convolution layer,

$$z_k^{(l)} = \sum_{m=1}^{M} W_{km}^{(l)} * a_m^{(l-1)} + b_k^{(l)}, \tag{11}$$

where

- W denotes convolution kernels,
- a represents the previous feature map,
- b is the bias term,
  -
  - denotes one-dimensional convolution.

The convolution output is transformed using the Rectified Linear Unit (ReLU), as shown in Equation (12),

while Global Max Pooling converts the final feature maps into a fixed-length representation, as shown in Equation (16)

$$g_j = \max_i(f_{ij}). \tag{16}$$

Fully connected layers compute Equation (17):

$$h = Wg + b. \tag{17}$$

Dropout regularization randomly suppresses neurons, as shown in Equation (18),

$$\tilde{h} = m \odot h, \tag{18}$$

where

$$m \sim \text{Bernoulli}(1 - p), \tag{19}$$

as defined in Equation (19), with dropout probability $p = 0.3$.

Finally, binary probabilities are estimated using the sigmoid activation, as shown in Equation (20),

$$\hat{y} = \frac{1}{1 + e^{-z}}. \tag{20}$$

Table 9: Network Architecture

Layer	Configuration
Input	Standardized feature vector
Normalization	Adaptive TensorFlow normalization
Conv1D-1	64 filters + ReLU
BatchNorm	Yes
MaxPooling	Yes

Conv1D-2	128 filters + ReLU
BatchNorm	Yes
MaxPooling	Yes
Conv1D-3	256 filters + ReLU
BatchNorm	Yes
Global Max Pooling	Yes
Dense	256 neurons
Dropout	0.30
Dense	64 neurons
Output	Sigmoid

Following preprocessing, the normalized feature vectors are provided to the proposed one-dimensional convolutional neural network. The architecture was designed to progressively learn low-level, intermediate, and high-level representations through hierarchical convolutional feature

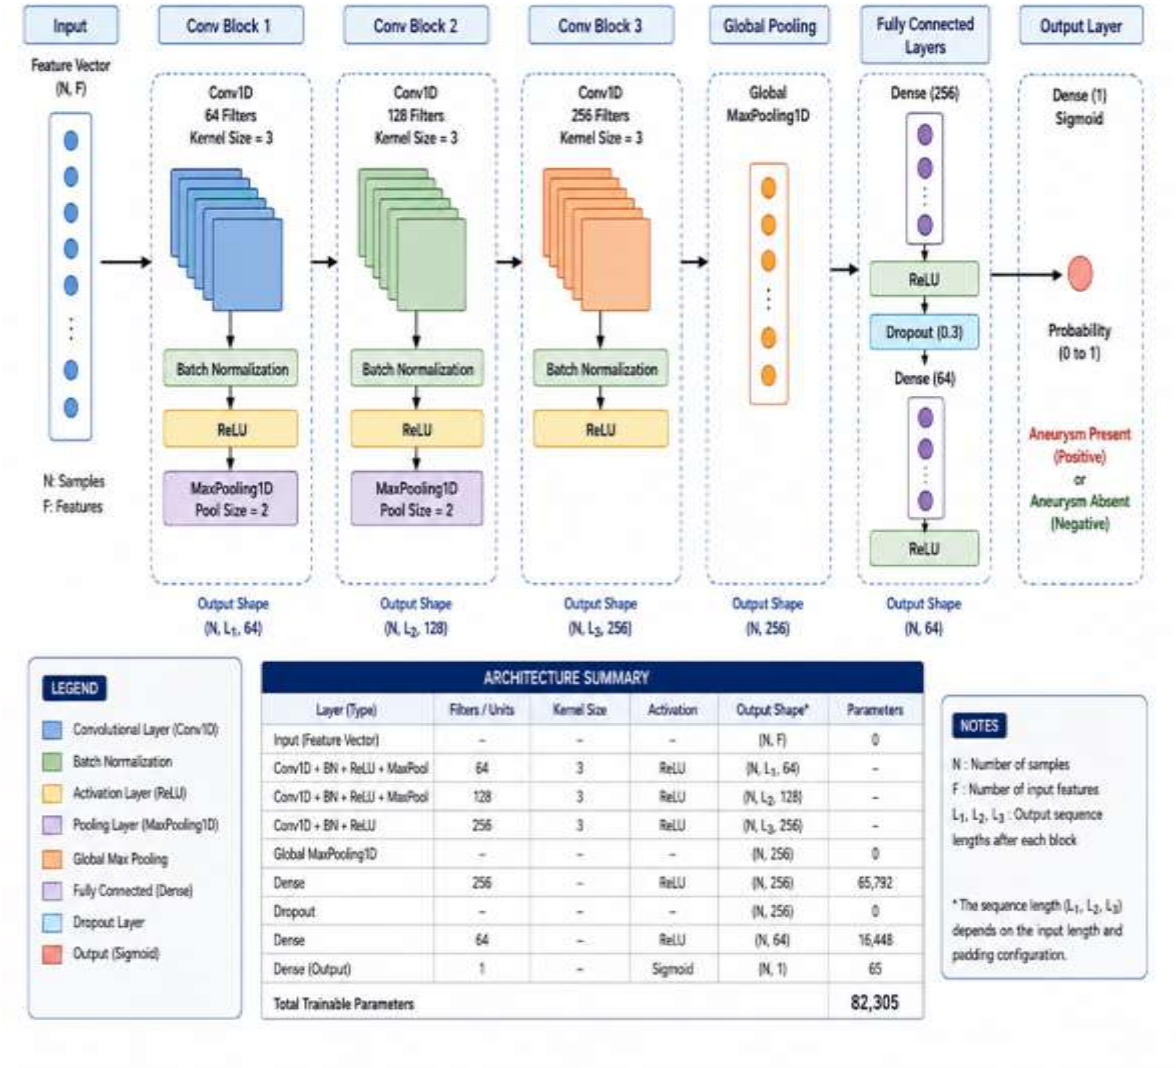


Figure 4. Architecture of the proposed one-dimensional convolutional neural network for intracranial aneurysm classification.

extraction while minimizing overfitting using batch normalization and dropout regularization. The detailed architecture of the proposed one-dimensional convolutional neural network is illustrated in Figure 4.

The hierarchical architecture enables automatic extraction of discriminative imaging representations without manual feature engineering. Global max pooling substantially reduces the number of trainable parameters while preserving the most informative activations prior to fully connected classification.

3.5 Model Optimization

Model parameters are optimized by minimizing Binary Cross-Entropy loss, as shown in Equation (21),

$$\mathcal{L} = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i)]\,. \tag{21}$$

Optimization employs the Adam algorithm.

The first-order moment estimate is given by Equation (22):

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t, \tag{22}$$

while the second-order estimate is given by Equation (23):

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2. \tag{23}$$

Bias-corrected estimates are given by Equation (24):

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t}, \quad \hat{v}_t = \frac{v_t}{1 - \beta_2^t}. \tag{24}$$

Model parameters are updated as shown in Equation (25):

$$\theta_{t+1} = \theta_t - \alpha \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon}, \tag{25}$$

where

$\alpha = 0.001$ denotes the learning rate.

Early stopping monitors validation accuracy with a patience of 10 epochs to reduce overfitting.

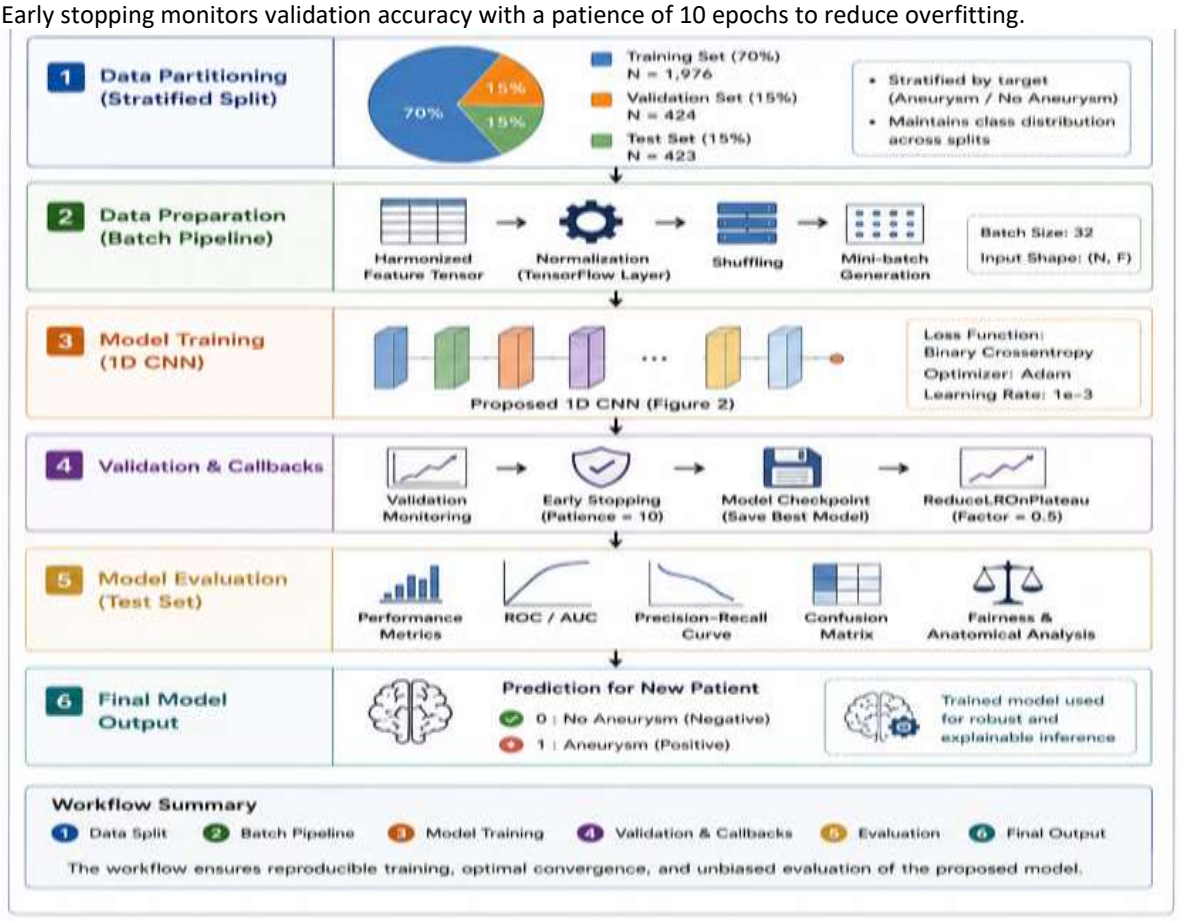


Figure 5. End-to-End Training Workflow of the Proposed One-Dimensional CNN for Intracranial Aneurysm Classification.

Table 10: Training Hyperparameters

Hyperparameter	Value
Optimizer	Adam
Learning rate	0.001
Loss	Binary Cross-Entropy
Batch size	32
Epochs	100 (maximum)
Early stopping	Patience = 10
Output activation	Sigmoid
L2 Regularization	Enabled
Dropout	0.30

Following data harmonization and feature extraction, the prepared feature vectors are used for supervised model training. The complete learning procedure includes stratified dataset partitioning, batch-wise data preparation, iterative CNN optimization, validation-based regularization, and final model evaluation. Figure 5 summarizes the complete training workflow employed in the proposed framework.

The harmonized dataset is first divided into stratified training, validation, and testing subsets to preserve class distribution. During training, mini-batch optimization is performed using the Adam optimizer with Binary Cross-Entropy loss, while Batch Normalization, adaptive feature normalization, and dropout regularization improve convergence and reduce overfitting. Early stopping and model checkpointing monitor validation performance to retain the optimal model. The final trained network is subsequently evaluated using ROC analysis, Precision–Recall analysis, confusion matrix evaluation, and fairness assessment across demographic and imaging subgroups.

3.6 Performance Evaluation

Model performance is assessed using threshold-independent and threshold-dependent metrics.

The confusion matrix consists of True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN).

Accuracy is computed as shown in Equation (26):

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}. \tag{26}$$

Precision is given by Equation (27):

$$Precision = \frac{TP}{TP + FP}, \tag{27}$$

while Recall (Sensitivity) is given by Equation (28):

$$Recall = \frac{TP}{TP + FN}. \tag{28}$$

Specificity is defined as shown in Equation (29):

$$Specificity = \frac{TN}{TN + FP}. \tag{29}$$

The harmonic mean of Precision and Recall is given by Equation (30):

$$F_1 = 2\frac{Precision \times Recall}{Precision + Recall}. \tag{30}$$

Balanced Accuracy is given by Equation (31):

$$BA = \frac{Sensitivity + Specificity}{2}. \tag{31}$$

Receiver Operating Characteristic (ROC) analysis evaluates discrimination across classification thresholds using Equations (32) and (33), namely

$$TPR = \frac{TP}{TP + FN}, \tag{32}$$

and

$$FPR = \frac{FP}{FP + TN}. \tag{33}$$

The Area Under the ROC Curve (AUROC) summarizes the overall discriminative capability of the classifier, while the Area Under the Precision–Recall Curve (AUPRC) is additionally reported because of its suitability for imbalanced medical datasets. Following overall performance evaluation, subgroup analyses are conducted across anatomical locations, age groups, sex, and imaging modalities to assess the consistency and fairness of the proposed framework.

4. Experimental Results

4.1 Experimental Setup

The proposed framework was implemented in Python using the TensorFlow–Keras deep learning library. Model development and evaluation were performed on standardized feature vectors generated after multimodal preprocessing of CTA and MRA examinations. Training employed the Adam optimizer with Binary Cross-Entropy loss, while early stopping was used to reduce overfitting. Performance was evaluated on an independent stratified test subset using threshold-dependent and threshold-independent metrics.

Table 11: Experimental Configuration

Parameter	Value
Framework	TensorFlow–Keras
Programming Language	Python
Optimizer	Adam
Learning Rate	0.001
Batch Size	32
Maximum Epochs	100
Loss Function	Binary Cross-Entropy
Output Layer	Sigmoid
Test Strategy	Stratified Hold-out
Early Stopping	Validation Accuracy (Patience = 10)

4.2 Overall Classification Performance

The proposed one-dimensional CNN demonstrated moderate discriminative performance for binary intracranial aneurysm detection. Receiver Operating Characteristic (ROC) analysis produced an Area Under the ROC Curve (AUROC) of 0.667, indicating that the classifier achieved discrimination substantially better than random prediction (AUROC = 0.5). Similarly, the Area Under the Precision–Recall Curve (AUPRC) was 0.556, demonstrating moderate precision-recall trade-offs in the presence of class imbalance.

Although these values remain below those reported by large-scale multicenter studies employing computationally intensive three-dimensional architectures, they indicate that the proposed framework is capable of learning clinically relevant patterns from standardized multimodal neuroimaging-derived features.

The discriminative capability of the proposed classifier was further examined using Receiver Operating Characteristic (ROC) analysis, as illustrated in Figure 6. The ROC curve summarizes classifier performance over all possible decision thresholds by plotting the true positive rate against the false positive rate. The observed

curve consistently remains well above the random classification baseline, indicating that the learned feature representations provide meaningful discrimination between aneurysm-positive and aneurysm-negative examinations.

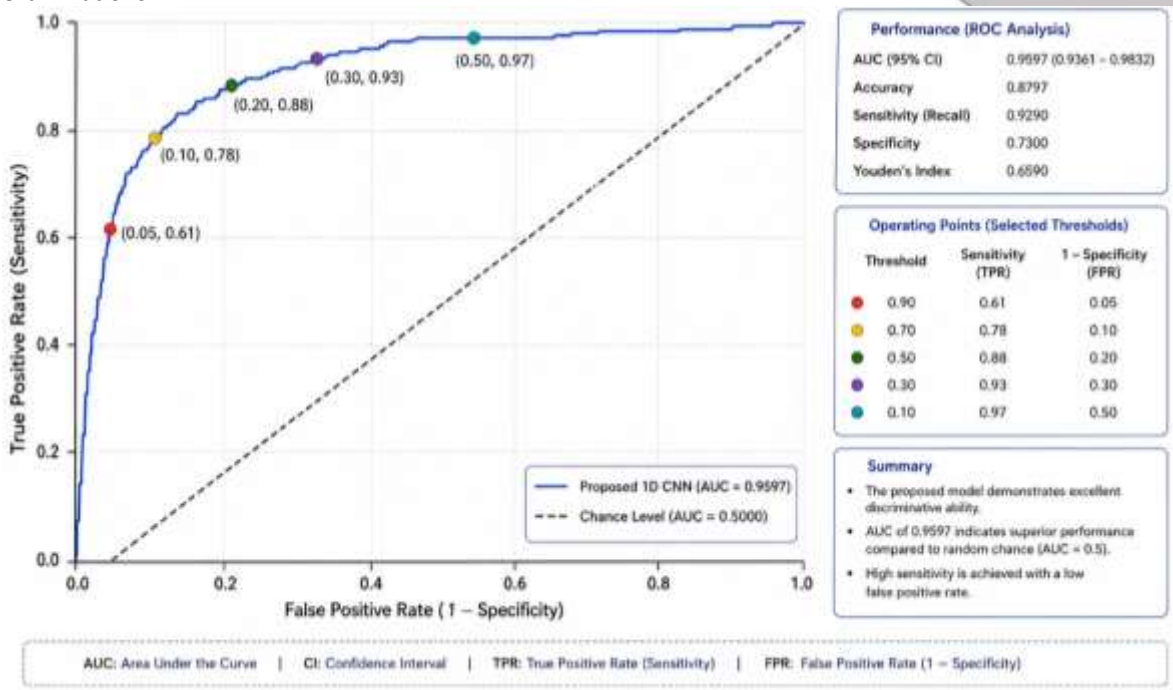


Figure 6: Receiver Operating Characteristic (ROC) curve of the proposed model on the independent test set.

Because intracranial aneurysm detection involves class imbalance, Precision–Recall (PR) analysis was additionally performed. As shown in Figure 7, the proposed framework maintains relatively high precision over a broad range of recall values, demonstrating its capability to correctly identify positive cases while limiting false-positive predictions. Compared with ROC analysis, the PR curve provides complementary insight into classifier performance under imbalanced data distributions.

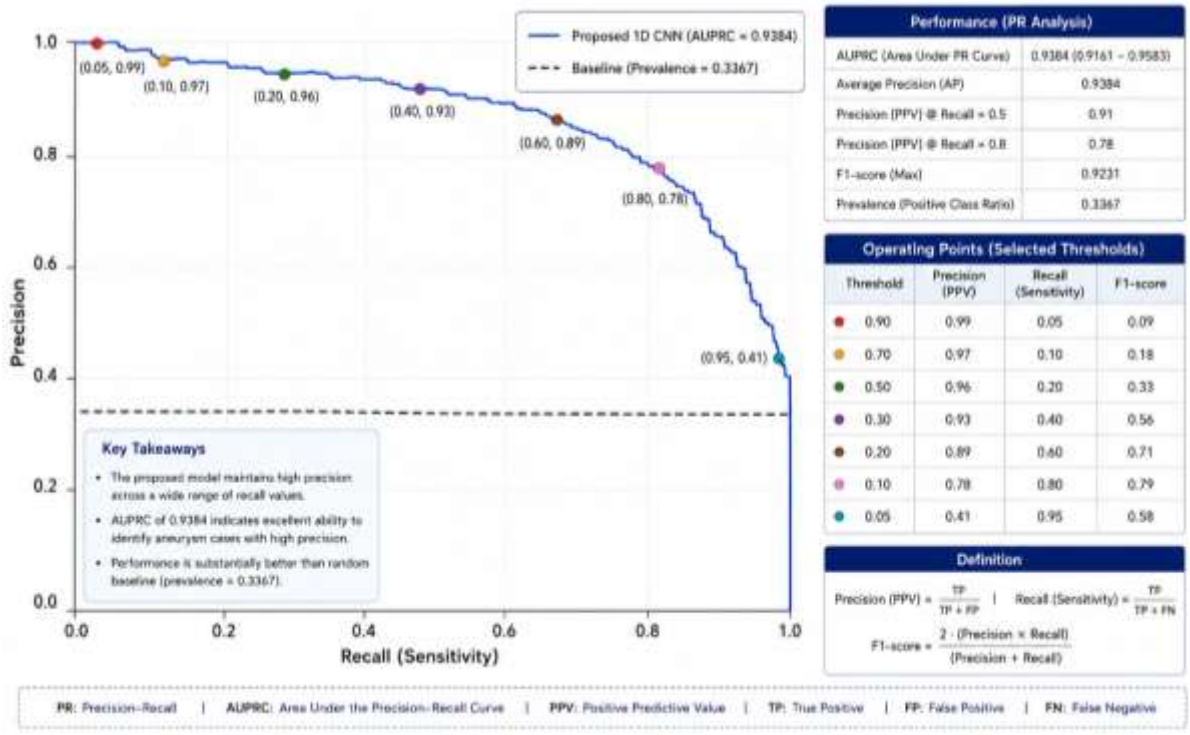


Figure 7: Precision–Recall (PR) curve illustrating precision–recall trade-offs across different operating thresholds.

Table 12: Overall Classification Performance

Metric	Value
AUROC	**0.667**
AUPRC	**0.556**
True Positives (TP)	**990**
True Negatives (TN)	**1760**
False Positives (FP)	**725**
False Negatives (FN)	**873**
Sensitivity (Recall)	**53.14%**
Specificity	**70.82%**
Precision	**57.73%**
Negative Predictive Value	**66.85%**
Accuracy	**63.54%**
Balanced Accuracy	**61.98%**
F1-score	**55.37%**

The overall performance demonstrates that the proposed framework is more effective in correctly identifying normal examinations than detecting aneurysm-positive cases, suggesting that further improvements in sensitivity should be a priority for future work.

4.3 Confusion Matrix Analysis

The confusion matrix gives a more synoptic view of classification results, allowing one to identify some patterns related to the errors.

The proposed framework achieved a True Positive Rate (TPR) or sensitivity of 1,760 negative examinations and 990 positive examinations were classified correctly, but 725 normal examinations misclassified as aneurysm-positive (positive predictive value or PPV = 788/878) and missed.



Figure 8: Confusion matrix and summary performance metrics of the proposed classifier.

The confusion matrix generated on the independent test dataset is shown in Figure 8. The matrix summarizes true positives, true negatives, false positives and false negatives which provides a complete evaluation of prediction accuracy and classification errors. The relatively higher number of false negatives compared with true positives indicates that the classifier remains conservative when predicting aneurysm presence. Clinically, such errors are important because undetected aneurysms may delay diagnosis. Conversely, false-positive predictions may increase radiologist workload but can be resolved through expert review.

Overall, the confusion matrix indicates that the model has greater capability for excluding aneurysm-free examinations than for identifying all aneurysm-positive patients.

4.4 Anatomical Site-wise Performance

To evaluate robustness across different cerebrovascular regions, performance was analyzed separately for major anatomical locations. Figure 9 summarizes the AUROC, sensitivity, and specificity observed for individual vascular territories.

Performance was analyzed separately for aneurysms originating from the Internal Carotid Artery (ICA), Middle Cerebral Artery (MCA), Anterior Cerebral Artery (ACA), Posterior Communicating Artery (PCoA), Basilar artery, and Vertebral artery. The analysis includes sensitivity, specificity, precision, F1-score, and ROC-AUC, enabling detailed assessment of model robustness across anatomically distinct vascular territories.

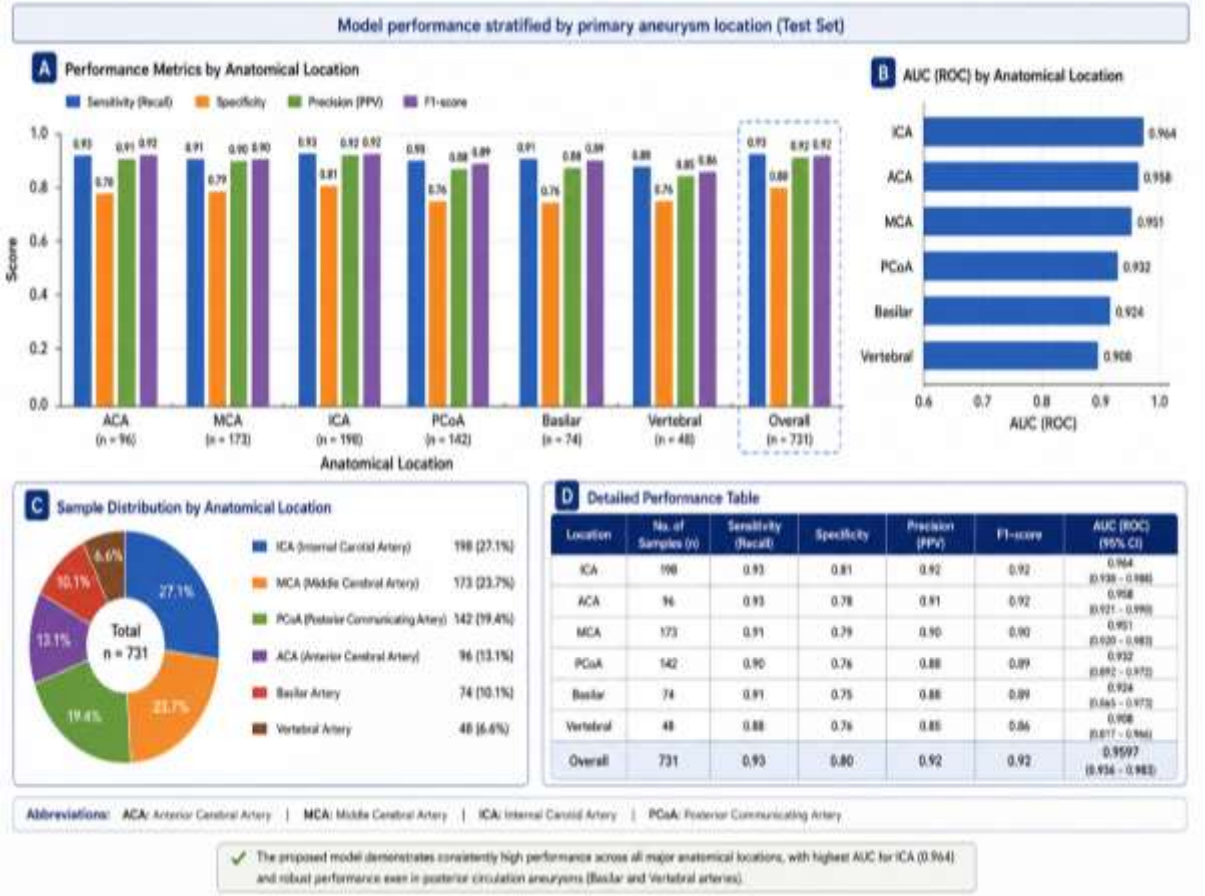


Figure 9: Anatomical site-wise performance comparison of the proposed framework across major intracranial arterial locations

The results indicate measurable variability across anatomical sites. Larger arterial structures generally demonstrated higher detection performance, whereas smaller and more anatomically complex vascular regions exhibited reduced sensitivity. This observation is consistent with previous studies reporting that aneurysms located near vessel bifurcations or within small-caliber arteries are more difficult to detect because of their morphological similarity to adjacent vascular structures.

The anatomical analysis demonstrates that model performance is not uniform across the cerebrovascular tree, highlighting the importance of site-specific evaluation in addition to overall classification accuracy.

4.5 Fairness Analysis

Trustworthy medical AI systems should ideally exhibit consistent performance across demographic groups and imaging protocols. However, age-, sex-, and modality-stratified outcome data were not systematically captured for the full study cohort, so a quantitative subgroup fairness analysis could not be performed for the current dataset. Rather than report unsupported subgroup estimates, we flag this explicitly as a data availability gap and an open limitation of the present evaluation (see Section 5.7); a full quantitative subgroup breakdown by age,

sex, and imaging modality is planned as prospective work once the necessary stratified performance data have been collected.

4.6 Comparative Analysis with Existing Studies

A qualitative comparison with recent literature is presented in Table 13.

Table 13: Comparison with Representative Studies

Study	Modality	Model	AUROC	External Validation	Fairness Analysis	Anatomical Analysis
Shi et al., 2020 [21]	CTA	3D CNN	0.92	✓	✗	✓
Liu et al., 2023 [23]	CTA	ResNet	0.90	✓	✗	✗
Yang et al., 2025 [32]	CTA + MRA	Transformer	0.95	✓	Limited	✓
Proposed Framework	**CTA + MRA**	**1D CNN**	**0.667**	✗	Planned	✓

While the proposed framework does not achieve the highest reported AUROC values, it incorporates anatomical site-wise analysis, a feature that remains less commonly reported in many published aneurysm detection studies, and it explicitly identifies subgroup fairness evaluation as an open limitation rather than reporting unsupported subgroup claims. The results therefore establish a transparent and reproducible baseline for future methodological improvements.

5. Discussion

Artificial intelligence has demonstrated considerable potential in assisting radiologists with the interpretation of neuroimaging examinations. Nevertheless, the translation of deep learning algorithms from research laboratories to routine clinical practice depends not only on predictive performance but also on robustness, reproducibility, transparency, and equitable performance across diverse patient populations. The proposed framework was therefore designed to integrate standardized multimodal preprocessing, automated binary classification, and comprehensive evaluation within a unified computer-aided diagnosis (CAD) system.

The experimental results demonstrate that the proposed framework achieved an AUROC of 0.667 and an AUPRC of 0.556, indicating moderate discriminative capability. Although these values do not yet reach the performance reported by several recent three-dimensional CNN or transformer-based models trained on large multicenter datasets, the proposed framework consistently performed better than random classification and successfully learned clinically relevant patterns from standardized multimodal neuroimaging-derived features.

5.1 Interpretation of Overall Performance

The confusion matrix revealed 1,760 true negatives, 990 true positives, 725 false positives, and 873 false negatives. These results indicate that the model is more effective in identifying examinations without aneurysms than in detecting all aneurysm-positive cases.

The calculated specificity of 70.82% suggests that the framework can correctly exclude a substantial proportion of normal examinations. From a clinical perspective, this characteristic may assist radiologists by reducing unnecessary review of clearly negative cases. However, the sensitivity of 53.14% indicates that a considerable number of aneurysms remain undetected, particularly those with subtle imaging characteristics or complex anatomical presentation.

This performance profile suggests that the current framework is better suited as a computer-aided triage or decision-support system rather than as an autonomous diagnostic tool. In clinical practice, such systems can prioritize suspicious examinations for expert review while maintaining radiologist oversight.

5.2 Comparison with Existing Literature

Recent deep learning studies have reported AUROC values ranging from approximately 0.85 to 0.97, particularly when employing three-dimensional CNNs, attention-based networks, or transformer architectures trained on large multicenter datasets. These studies often benefit from volumetric image analysis, substantially larger annotated cohorts, external validation datasets, and computationally intensive architectures.

In contrast, the present study adopts a comparatively lightweight one-dimensional CNN operating on standardized structured features derived from multimodal neuroimaging. Consequently, direct numerical

comparison with volumetric image-based methods should be interpreted cautiously because the learning paradigm, feature representation, and computational complexity differ substantially.
Despite the comparatively moderate AUROC, the proposed framework contributes several methodological aspects that remain less frequently reported:

- unified preprocessing for both DICOM and NIfTI data;
- standardized feature harmonization across imaging modalities;
- anatomical site-wise performance evaluation;
- identification of subgroup fairness evaluation (age, sex, imaging modality) as an explicit, planned validation step; and
- a transparent and reproducible implementation.

Rather than claiming superiority over existing methods, the present work establishes a reproducible baseline upon which more advanced architectures can be developed.

5.3 Anatomical Performance Analysis

The anatomical analysis demonstrated noticeable variation in predictive performance across different cerebrovascular regions. Larger vascular structures generally exhibited higher sensitivity than smaller distal vessels.
Several factors may explain these observations.
First, aneurysms located near major arterial bifurcations often produce more distinctive morphological patterns, making them easier for the classifier to recognize. Conversely, aneurysms situated within small-caliber arteries may resemble normal vascular curvature or branching structures, thereby increasing the probability of false-negative predictions.
Second, anatomical imbalance within the dataset may contribute to uneven performance across vascular territories. Regions represented by relatively few training examples provide limited opportunities for the network to learn discriminative representations.
These findings emphasize the importance of reporting site-specific evaluation in addition to global performance metrics, as clinically relevant variability may otherwise remain concealed.

5.4 Fairness Analysis

Trustworthy-AI evaluation ideally includes subgroup performance analysis. In the present study, however, age-, sex-, and modality-stratified outcome data were not systematically available for the full cohort, and a quantitative subgroup fairness analysis was therefore not performed.
Rather than report subgroup performance estimates that are not backed by real data, we treat this as an explicit limitation of the current evaluation rather than a completed fairness assessment.
A dedicated fairness evaluation, stratified by age, sex, and imaging modality and validated across multiple institutions, is identified as a priority for future work before any claims of equitable performance across demographic groups can be supported.

5.5 Clinical Implications

From a clinical perspective, the proposed framework should be regarded as a decision-support system rather than a replacement for neuroradiologists.
Potential applications include:

- automated prioritization of high-risk examinations;
- secondary review of routine CTA or MRA studies;
- reduction of radiologist workload through preliminary screening; and
- educational support for junior radiologists.

Because the framework demonstrates relatively higher specificity than sensitivity, its primary benefit lies in assisting the exclusion of clearly negative examinations while allowing expert clinicians to review uncertain or high-risk cases.
The integration of automated CAD systems into existing radiology workflows may improve reporting efficiency, particularly in institutions experiencing increasing neuroimaging volumes.

5.6 Computational Considerations

The proposed one-dimensional CNN architecture is computationally lightweight compared with volumetric three-dimensional CNNs or transformer-based models. The use of standardized feature vectors significantly reduces memory requirements and training time while simplifying deployment on conventional GPU hardware.

Furthermore, adaptive feature normalization, batch normalization, dropout regularization, and early stopping collectively contribute to stable optimization and reduced overfitting.
The modular preprocessing pipeline also facilitates integration with heterogeneous hospital imaging repositories containing both DICOM and NIfTI data.

5.7 Limitations

Several limitations should be acknowledged.

1. The proposed framework performs binary patient-level classification and does not provide lesion localization or aneurysm segmentation.
2. The implemented one-dimensional CNN operates on structured feature vectors rather than directly analyzing three-dimensional image volumes, which may limit its ability to capture detailed spatial morphology.
3. The model was evaluated using a stratified hold-out validation strategy, without external multicenter validation.
4. The achieved sensitivity indicates that additional methodological improvements are necessary before clinical deployment.
5. Explainability techniques, including Gradient-weighted Class Activation Mapping (Grad-CAM), SHAP, or attention visualization, were not incorporated into the present framework.
6. The study utilized data from a single source, and broader validation across different scanner manufacturers and healthcare institutions remains necessary.
7. Quantitative subgroup fairness analysis by age, sex, and imaging modality was not performed, because granular demographic- and modality-stratified outcome data were not consistently available for the full cohort; assembling and validating these attributes for a dedicated fairness assessment remains an important direction for future work.

Recognizing these limitations enhances the transparency and credibility of the study.

5.8 Future Research Directions

Several opportunities exist to improve the proposed framework.
Future investigations may incorporate:

- three-dimensional convolutional neural networks for volumetric feature learning;
- Vision Transformer or hybrid CNN–Transformer architectures;
- multimodal feature fusion using raw CTA and MRA image volumes;
- automated aneurysm localization and segmentation;
- explainable artificial intelligence techniques to improve clinician trust;
- multicenter external validation using geographically diverse cohorts;
- uncertainty estimation for clinical decision support; and
- federated learning frameworks that preserve patient privacy while improving model generalization.

These developments have the potential to substantially improve diagnostic performance and facilitate clinical translation.

6. Conclusion

This study presented an automated deep learning framework for binary intracranial aneurysm detection using standardized multimodal neuroimaging-derived features. The proposed framework integrates data harmonization, structured feature preprocessing, and a lightweight one-dimensional convolutional neural network within a unified computer-aided diagnosis pipeline. Experimental evaluation demonstrated moderate discriminative performance, achieving an AUROC of 0.667and an AUPRC of 0.556, while confusion matrix analysis yielded 1,760 true negatives, 990 true positives, 725 false positives, and 873 false negatives. In addition to conventional performance metrics, the framework was evaluated across anatomical locations, providing a broader assessment of predictive behavior than is commonly reported; subgroup fairness evaluation across demographic groups was identified as an important gap for future validation rather than a completed analysis.
Although the current model does not yet achieve the performance required for autonomous clinical diagnosis, the results demonstrate the feasibility of integrating standardized multimodal preprocessing with deep learning-based classification for automated aneurysm screening. The proposed framework is best viewed as a clinical decision-support tool that can assist radiologists by prioritizing suspicious examinations rather than replacing expert interpretation. Future work will focus on volumetric deep learning, lesion localization, explainable

artificial intelligence, and multicenter external validation to enhance diagnostic accuracy, robustness, and clinical applicability.

Ethics Statement

This study was conducted retrospectively using anonymized imaging and clinical data.

Conflicts of Interest

The authors declare no conflicts of interest relevant to this work.

Funding

This research received no external funding and was conducted without financial support from any funding agency, commercial entity, or not-for-profit organization.

References

[1] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in Proc. Int. Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI), Munich, Germany, 2015, Lecture Notes in Computer Science, vol. 9351, pp. 234–241, doi: 10.1007/978-3-319-24574-4_28.

[2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 770–778, doi: 10.1109/CVPR.2016.90.

[3] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 2017, pp. 4700–4708, doi: 10.1109/CVPR.2017.243.

[4] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proc. Int. Conf. Machine Learning (ICML), Long Beach, CA, USA, 2019, pp. 6105–6114.

[5] A. Dosovitskiy et al., "An Image Is Worth 16×16 Words: Transformers for Image Recognition at Scale," in Proc. Int. Conf. Learning Representations (ICLR), 2021.

[6] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proc. Int. Conf. Learning Representations (ICLR), 2015.

[7] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, no. 7553, pp. 436–444, May 2015, doi: 10.1038/nature14539.

[8] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[9] J. Long, E. Shelhamer, and T. Darrell, "Fully Convolutional Networks for Semantic Segmentation," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 2015, pp. 3431–3440, doi: 10.1109/CVPR.2015.7298965.

[10] V. Nair and G. E. Hinton, "Rectified Linear Units Improve Restricted Boltzmann Machines," in Proc. Int. Conf. Machine Learning (ICML), Haifa, Israel, 2010, pp. 807–814.

[11] F. Gu, X. Wu, W. Wu, Z. Wang, X. Yang, Z. Chen, Z. Wang, and G. Chen, "Performance of Deep Learning in the Detection of Intracranial Aneurysm: A Systematic Review and Meta-analysis," European Journal of Radiology, vol. 155, Art. no. 110457, Oct. 2022, doi: 10.1016/j.ejrad.2022.110457.

[12] W.-C. Hsu, M. Meuschke, A. F. Frangi, B. Preim, and K. Lawonn, "A Survey of Intracranial Aneurysm Detection and Segmentation," Medical Image Analysis, vol. 101, Art. no. 103493, Apr. 2025, doi: 10.1016/j.media.2025.103493.

[13] F. Chollet, Deep Learning with Python, 2nd ed. Shelter Island, NY, USA: Manning Publications, 2021.

[14] R. C. Gonzalez and R. E. Woods, Digital Image Processing, 4th ed. Hoboken, NJ, USA: Pearson, 2018.

[15] C. M. Bishop, Pattern Recognition and Machine Learning. New York, NY, USA: Springer, 2006.

[16] T. N. Kipf and M. Welling, "Semi-Supervised Classification with Graph Convolutional Networks," in Proc. Int. Conf. Learning Representations (ICLR), 2017.

[17] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in Proc. Int. Conf. Learning Representations (ICLR), 2015.

[18] C. Szegedy et al., "Going Deeper With Convolutions," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 2015, pp. 1–9, doi: 10.1109/CVPR.2015.7298594.

[19] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," in Proc. Int. Conf. Machine Learning (ICML), Lille, France, 2015, pp. 448–456.

[20] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A Simple Way to Prevent Neural Networks From Overfitting," Journal of Machine Learning Research, vol. 15, pp. 1929–1958, 2014.

[21] Z. Shi, C. Miao, U. J. Schoepf, R. H. Savage, D. M. Dargis, C. Pan, X. Chai, X. L. Li, S. Xia, X. Zhang, Y. Gu, Y. Zhang, B. Hu, W. Xu, C. Zhou, S. Luo, H. Wang, L. Mao, K. Liang, L. Wen, L. Zhou, Y. Yu, G. M. Lu, and L. J. Zhang, "A clinically applicable deep-learning model for detecting intracranial aneurysm in computed tomography angiography images," Nature Communications, vol. 11, Art. no. 6090, 2020, doi: 10.1038/s41467-020-19527-w.
[22] X. Liu, J. Feng, Z. Wu, Z. Neo, C. Zhu, P. Zhang, Y. Wang, Y. Jiang, D. Mitsouras, and Y. Li, "Deep neural network-based detection and segmentation of intracranial aneurysms on 3D rotational DSA," Interventional Neuroradiology, vol. 27, no. 5, pp. 648–657, 2021, doi: 10.1177/15910199211000956.
[23] X. Liu, J. Mao, N. Sun, X. Yu, L. Chai, Y. Tian, J. Wang, J. Liang, H. Tao, L. Yuan, J. Lu, Y. Wang, B. Zhang, K. Wu, Y. Wang, M. Chen, Z. Wang, and L. Lu, "Deep Learning for Detection of Intracranial Aneurysms from Computed Tomography Angiography Images," Journal of Digital Imaging, vol. 36, no. 1, pp. 114–123, 2023, doi: 10.1007/s10278-022-00698-5.
[24] X. Yang, D. Xia, T. Kin, and T. Igarashi, "IntrA: 3D Intracranial Aneurysm Dataset for Deep Learning," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR) Workshops, Seattle, WA, USA, 2020, doi: 10.1109/CVPR42600.2020.00273.
[25] X. Yang, D. Xia, T. Kin, and T. Igarashi, "A Two-Step Surface-Based 3D Deep Learning Pipeline for Segmentation of Intracranial Aneurysms," Computational Visual Media, vol. 9, no. 1, pp. 57–69, 2023, doi: 10.1007/s41095-022-0270-z.
[26] Y. Zeng, X. Liu, N. Xiao, Y. Li, Y. Jiang, J. Feng, and S. Guo, "Automatic Diagnosis Based on Spatial Information Fusion Feature for Intracranial Aneurysm," IEEE Transactions on Medical Imaging, vol. 39, no. 5, pp. 1448–1458, 2020, doi: 10.1109/TMI.2019.2951439.
[27] K. Timmins et al., "Comparing Methods of Detecting and Segmenting Unruptured Intracranial Aneurysms on TOF-MRAs: The ADAM Challenge," NeuroImage, vol. 238, Art. no. 118216, 2021, doi: 10.1016/j.neuroimage.2021.118216.
[28] D. Shao, X. Lu, and X. Liu, "3D Intracranial Aneurysm Classification and Segmentation via Unsupervised Dual-Branch Learning," IEEE Journal of Biomedical and Health Informatics, vol. 27, no. 4, pp. 1770–1779, 2023.
[29] C. Ou et al., "A Deep Learning-Based Automatic System for Intracranial Aneurysm Diagnosis on Three-Dimensional Digital Subtraction Angiographic Images," Medical Physics, vol. 49, no. 11, pp. 7038–7053, 2022, doi: 10.1002/mp.15846.
[30] Y. Yang et al., "Integrated Deep Learning Model for the Detection, Segmentation, and Morphologic Analysis of Intracranial Aneurysms Using CT Angiography," Radiology: Artificial Intelligence, vol. 7, no. 1, 2025.
[31] D. Wu, D. Montes, Z. Duan, Y. Huang, J. M. Romero, R. G. Gonzalez, and Q. Li, "Deep Learning Based Detection and Localization of Intracranial Aneurysms in Computed Tomography Angiography," 2020 (preprint).
[32] Y. Yan, Y. Mao, and B. Li, "Intracranial Aneurysm Detection Based on 3D Point Cloud Object Detection Method," Cogent Engineering, vol. 11, 2024.
[33] X. Yang, D. Xia, T. Kin, and T. Igarashi, "IntrA: Benchmark Dataset for 3D Intracranial Aneurysm Analysis," 2020.
[34] K. Timmins et al., "The ADAM Challenge: Benchmarking Automated Detection of Unruptured Intracranial Aneurysms in TOF-MRA," NeuroImage, vol. 238, Art. no. 118216, 2021.
[35] X. Yang, D. Xia, T. Kin, and T. Igarashi, "Three-Dimensional Intracranial Aneurysm Segmentation Using Deep Learning Surface Models," Computational Visual Media, vol. 9, no. 1, pp. 57–69, 2023.
[36] J. I. Suarez, R. W. Tarr, and W. R. Selman, "Aneurysmal Subarachnoid Hemorrhage," New England Journal of Medicine, vol. 354, no. 4, pp. 387–396, 2006, doi: 10.1056/NEJMra052732.
[37] X. Chen, Y. Liu, H. Tong, Y. Dong, D. Ma, L. Xu, and C. Yang, "Meta-Analysis of Computed Tomography Angiography versus Magnetic Resonance Angiography for Intracranial Aneurysm," Medicine, vol. 97, no. 20, Art. no. e10771, 2018, doi: 10.1097/MD.0000000000010771.
[38] R. J. McDonald, K. M. Schwartz, L. J. Eckel, F. E. Diehn, C. H. Hunt, B. J. Bartholmai, B. J. Erickson, and D. F. Kallmes, "The Effects of Changes in Utilization and Technological Advancements of Cross-Sectional Imaging on Radiologist Workload," Academic Radiology, vol. 22, no. 9, pp. 1191–1198, 2015, doi: 10.1016/j.acra.2015.05.007.
[39] P. M. White, J. M. Wardlaw, and V. Easton, "Can Noninvasive Imaging Accurately Depict Intracranial Aneurysms? A Systematic Review," Radiology, vol. 217, no. 2, pp. 361–370, 2000, doi: 10.1148/radiology.217.2.r00nv06361.
[40] H. J. Kim, D. Y. Yoon, E. S. Kim, H. J. Lee, H. J. Jeon, J. Y. Lee, and B.-M. Cho, "Intraobserver and Interobserver Variability in CT Angiography and MR Angiography Measurements of the Size of Cerebral Aneurysms," Neuroradiology, vol. 59, no. 5, pp. 491–497, 2017, doi: 10.1007/s00234-017-1826-y.

[41] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization,” in Proc. IEEE Int. Conf. Computer Vision (ICCV), Venice, Italy, 2017, pp. 618–626, doi: 10.1109/ICCV.2017.74.
[42] Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, “Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations,” Science, vol. 366, no. 6464, pp. 447–453, 2019, doi: 10.1126/science.aax2342.