Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 3, 2026

AN ADVANCED ATTENTION-ENHANCED CNN–BILSTM FRAMEWORK FOR MULTIMODAL BRAIN-IMAGING ANALYSIS, EARLY DIAGNOSIS, AND DISEASE-SEVERITY CLASSIFICATION IN PATIENTS WITH ISCHEMIC STROKE

Afaq Ahmad[*1], Ammara Arshad Javaid[2], Syed Ahmed Raza[3], Muhammad Sajjad[4]

*[*1]College of Artificial Intelligence (School of Software), Yanshan University, Qinhuangdao, 066004 China.*
[2]School of Information Science and Engineering, Communication Engineering, Yanshan University, Qinhuangdao, 066004 China
[3]Department of Computer Science, SZABIST University Hyderabad Campus, Pakistan
[4]School of Computing Technologies, RMIT University, Melbourne, VIC 3000, Australia

[*1]afaqkhan@stumail.ysu.edu.cn, [2]ammaraarshad@stumail.ysu.edu.cn, [3]ahmed.raza@hyd.szabist.edu.pk, [4]msajad32@gmail.com

DOI: https://doi.org/10.5281/zenodo.23914290

Keywords
Ischemic Stroke; Multimodal Brain Imaging; Medical Image Processing; Computer Vision; Attention-Enhanced CNN–BiLSTM; Deep Learning; Early Diagnosis; NIHSS-Based Severity Classification; Clinical Decision Support.

Article History
Received: 25 January 2026
Accepted: 10 March 2026
Published: 24 March 2026

Copyright @Author
Corresponding Author: *
Afaq Ahmad

Abstract

Ischemic stroke represents a major global healthcare challenge due to its substantial contribution to mortality, long-term neurological disability, and increasing clinical demands for rapid and accurate diagnosis. Despite advances in neuroimaging technologies, early identification of ischemic abnormalities and reliable assessment of neurological severity remain challenging because of heterogeneous lesion characteristics, subtle pathological changes, and limitations in conventional image interpretation. This study presents an advanced Attention-Enhanced Convolutional Neural Network–Bidirectional Long Short-Term Memory (CNN–BiLSTM) framework integrating medical image processing, computer vision, and multimodal deep learning for automated ischemic stroke diagnosis and disease-severity classification. The investigation utilizes a comprehensive dataset comprising 12,500 multimodal brain-imaging examinations obtained from 2,500 patients, including 1,800 ischemic stroke patients and 700 non-stroke controls. Four complementary neuroimaging modalities, namely computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC), are incorporated to capture diverse anatomical and pathological characteristics. The proposed framework employs image registration, intensity normalization, noise reduction, and data augmentation, followed by CNN-based spatial feature extraction, BiLSTM-based contextual dependency modeling, and attention-guided feature refinement. Furthermore, neurological impairment is categorized into mild, moderate, and severe stages using the National Institutes of Health Stroke Scale (NIHSS). Experimental results demonstrate superior diagnostic performance, achieving 97.8% accuracy, 97.2% precision, 96.9% sensitivity, 98.1% specificity, an F1-score of 97.0%, and an area under the receiver operating characteristic curve of 0.989. The proposed

architecture additionally attains 95.6% accuracy in multiclass disease-severity classification, outperforming conventional CNN, ResNet50, standalone BiLSTM, and CNN–LSTM architectures. Comparative and ablation analyses demonstrate the effectiveness of multimodal feature integration and attention-based representation learning in improving diagnostic discrimination and classification reliability. The findings establish the potential of the proposed CNN–BiLSTM framework to support accurate computer-assisted stroke identification, automated neurological severity assessment, and data-driven clinical decision-making. Overall, this research contributes a unified computational approach for multimodal neuroimaging interpretation, providing a promising foundation for intelligent stroke diagnosis and severity prediction in clinical healthcare environments, while highlighting the importance of further external validation.

1- Introduction:

Ischemic stroke is one of the most prevalent and life-threatening neurological disorders worldwide, contributing significantly to mortality, permanent neurological disability, and increasing healthcare expenditures. It occurs when cerebral blood circulation is disrupted because of arterial obstruction, resulting in insufficient oxygen and nutrient delivery to brain tissues and subsequently triggering neuronal dysfunction, cellular injury, and irreversible cerebral infarction. The severity and progression of ischemic stroke depend on several factors, including the location of vascular occlusion, duration of ischemia, cerebral collateral circulation, infarct volume, and extent of neurological impairment. Early diagnosis and accurate disease-severity assessment are essential for supporting timely therapeutic intervention, improving clinical decision-making, and minimizing long-term neurological complications [1]. However, identifying ischemic abnormalities during the early stages remains challenging because of subtle pathological changes, heterogeneous lesion characteristics, and considerable variations in individual patient presentations. Medical neuroimaging plays a fundamental role in the detection, characterization, and evaluation of ischemic stroke. Computed tomography (CT) and magnetic resonance imaging (MRI) are widely used to investigate structural abnormalities, identify cerebral lesions, and support diagnostic decision-making. CT imaging is particularly valuable in emergency healthcare environments because of its rapid acquisition, accessibility, and ability to detect intracranial hemorrhage. Nevertheless, conventional CT may demonstrate limited sensitivity for detecting subtle ischemic abnormalities during the earliest stages of cerebral infarction. Advanced MRI sequences, including diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC) maps, provide complementary information regarding cerebral tissue characteristics and pathological changes. DWI is particularly sensitive to restricted water diffusion associated with acute ischemic injury, whereas ADC maps assist in characterizing diffusion abnormalities [2]. FLAIR imaging provides additional information regarding tissue alterations and lesion evolution. The integration of these imaging modalities offers opportunities for more comprehensive analysis of cerebral ischemia and improved characterization of stroke-related abnormalities. Despite advancements in medical imaging technologies, conventional neuroimaging interpretation continues to face considerable clinical and computational limitations. Ischemic lesions frequently demonstrate variations in shape, size, anatomical location, intensity distribution, and pathological progression, making accurate identification particularly difficult in complex imaging environments. Furthermore, imaging artifacts, low contrast, scanner-dependent variations, and differences in acquisition protocols may reduce image consistency and diagnostic reliability. Traditional radiological assessment requires

extensive clinical expertise and may involve observer-dependent interpretation, particularly when examining subtle pathological abnormalities. These limitations have encouraged the development of advanced medical image processing, computer vision, and artificial intelligence techniques capable of automatically extracting meaningful imaging characteristics and supporting consistent neurological assessment. Recent developments in artificial intelligence (AI) and deep learning have introduced powerful computational approaches for automated medical image analysis and neurological disease classification [3]. Traditional machine learning methods commonly depend on manually engineered characteristics, including texture measurements, intensity distributions, morphological descriptors, and statistical features. Although these techniques have demonstrated potential in medical image classification, their effectiveness may be limited when analyzing high-dimensional neuroimaging datasets containing complex spatial relationships and heterogeneous pathological patterns. Deep learning architectures overcome several of these limitations through hierarchical representation learning, enabling computational models to automatically identify discriminative imaging features. Convolutional Neural Networks (CNNs), in particular, have become important components of medical computer vision systems because of their effectiveness in extracting localized spatial information, identifying anatomical structures, and recognizing complex pathological characteristics from multidimensional imaging data.

However, conventional CNN architectures primarily emphasize spatial feature extraction and may not adequately capture contextual relationships among consecutive brain-image slices or structured feature sequences. Ischemic stroke imaging involves complex anatomical dependencies that extend across multiple regions and imaging planes. Consequently, analyzing individual imaging slices independently may overlook important contextual information associated with lesion distribution and pathological progression. Long Short-Term Memory (LSTM) networks provide a mechanism for modeling dependencies within sequential representations, while Bidirectional Long Short-Term Memory (BiLSTM) networks extend this capability by processing feature sequences in both forward and backward directions. The integration of CNN and BiLSTM architectures therefore offers a promising strategy for combining detailed spatial representation learning with contextual information modeling. Such hybrid architectures can potentially improve the characterization of ischemic abnormalities by considering both localized imaging features and relationships among ordered anatomical representations. Attention mechanisms have further strengthened deep learning architectures by enabling adaptive identification and prioritization of informative feature representations. In ischemic stroke imaging, pathological abnormalities may occupy relatively small cerebral regions and exhibit imaging characteristics that are difficult to distinguish from surrounding healthy tissues [4]. Conventional feature aggregation approaches may therefore assign insufficient importance to subtle but diagnostically meaningful abnormalities. Attention-based learning addresses this limitation by assigning adaptive weights to feature representations according to their relevance to the prediction task. Integrating attention mechanisms with CNN–BiLSTM architectures may enhance discriminative feature learning, reduce the influence of less informative representations, and improve diagnostic classification. Furthermore, attention-guided feature refinement provides an additional computational mechanism for investigating which learned representations contribute to classification decisions, although attention weights alone do not establish clinical explainability. In addition to automated disease identification, accurate classification of neurological severity is an important component of comprehensive stroke assessment. The National Institutes of Health Stroke Scale (NIHSS) is a standardized clinical assessment instrument used to quantify neurological impairment by evaluating motor function, sensory response, consciousness, language, visual disturbances, and other neurological deficits. NIHSS-derived severity

categories provide a structured basis for investigating differences in clinical impairment among patients with ischemic stroke. However, neurological severity does not depend exclusively on lesion size or imaging characteristics, as functional deficits are also influenced by lesion location, affected neural pathways, and individual clinical conditions. Many automated stroke classification approaches primarily focus on distinguishing stroke patients from non-stroke controls, without integrating neurological severity assessment into the same computational architecture. Developing intelligent systems capable of supporting both ischemic stroke identification and NIHSS-associated severity classification may therefore contribute to more comprehensive computer-assisted neurological evaluation. To address these limitations, the present study proposes an advanced Attention-Enhanced Convolutional Neural Network–Bidirectional Long Short-Term Memory (CNN–BiLSTM) framework for multimodal brain-imaging analysis, early ischemic stroke diagnosis, and neurological severity classification [5]. The proposed architecture combines medical image processing, computer vision, and multimodal deep learning within a unified computational pipeline. The investigation utilizes a dataset comprising 12,500 multimodal brain-imaging examinations obtained from 2,500 patients, including 1,800 ischemic stroke patients and 700 non-stroke controls. Four complementary neuroimaging modalities, namely CT, DWI, FLAIR, and ADC, are incorporated to characterize anatomical and pathological variations associated with cerebral ischemia. Image registration, intensity normalization, noise reduction, and data augmentation are employed to prepare imaging data for subsequent analysis. CNN layers extract hierarchical spatial features, BiLSTM networks model contextual relationships among ordered feature representations, and attention mechanisms refine feature importance to strengthen diagnostic discrimination. The proposed framework investigates two interconnected classification objectives: automated ischemic stroke identification and neurological severity classification into mild, moderate, and severe categories using NIHSS-derived clinical labels. The reported experimental findings demonstrate a diagnostic accuracy of 97.8%, precision of 97.2%, sensitivity of 96.9%, specificity of 98.1%, F1-score of 97.0%, and area under the receiver operating characteristic curve (ROC–AUC) of 0.989. Additionally, the proposed architecture achieves 95.6% accuracy in multiclass neurological severity classification. Comparative evaluation against conventional CNN, ResNet50, standalone BiLSTM, and CNN–LSTM architectures is employed to assess the effectiveness of the proposed hybrid framework. Furthermore, ablation experiments investigate the contributions of multimodal feature integration, bidirectional contextual modeling, and attention-guided representation learning. These findings suggest that integrating complementary neuroimaging characteristics with advanced deep learning techniques may support more accurate and comprehensive computer-assisted stroke analysis.

The major contributions of this study are summarized as follows:

1. An advanced Attention-Enhanced CNN–BiLSTM architecture is proposed by integrating convolutional spatial feature extraction, bidirectional contextual dependency modeling, and adaptive attention mechanisms to support accurate automated ischemic stroke identification.
2. A comprehensive multimodal medical image processing framework is developed by combining CT, DWI, FLAIR, and ADC neuroimaging modalities with image registration, intensity normalization, noise reduction, and data augmentation to enhance pathological feature representation.
3. A unified deep learning approach is introduced to perform both ischemic stroke diagnosis and NIHSS-associated neurological severity classification across mild, moderate, and severe categories, supporting more comprehensive computational neurological assessment.
4. The proposed architecture is evaluated using 12,500 multimodal brain-imaging examinations from 2,500 patients, with reported performance of 97.8% diagnostic accuracy, 0.989 ROC–AUC,

and 95.6% multiclass severity classification accuracy.

This study contributes an integrated computational framework for multimodal neuroimaging interpretation, automated ischemic stroke diagnosis, and neurological severity assessment. By combining advanced image processing techniques with CNN-based spatial learning, BiLSTM-based contextual modeling, and attention-guided feature refinement, the proposed approach aims to address limitations associated with heterogeneous imaging characteristics and complex pathological representations. The investigation further supports research into intelligent computer-assisted neurological assessment and data-driven clinical decision-support systems. Nevertheless, additional evaluation using independent clinical datasets, different imaging acquisition environments, and prospective validation remains necessary to establish the generalizability and clinical applicability of the proposed framework. The remainder of this paper presents the related literature, research methodology, proposed model architecture, experimental evaluation, results and discussion, and concluding observations.

2- Medical Image Processing and Computer Vision for Brain Lesion Analysis:

Medical image processing and computer vision have become fundamental components of intelligent neurological diagnostic systems, providing advanced computational techniques for identifying, characterizing, and interpreting pathological abnormalities associated with ischemic stroke. The increasing availability of high-resolution medical imaging technologies has created opportunities for automated analysis of complex cerebral structures and disease-related abnormalities. However, conventional neuroimaging examination continues to encounter challenges associated with low tissue contrast, imaging artifacts, heterogeneous lesion morphology, and variations in image acquisition protocols. These limitations are particularly important during the early stages of ischemic stroke, when pathological alterations may be subtle and difficult to distinguish from normal brain tissues. Advanced medical image processing techniques address these challenges through image enhancement, spatial normalization, noise reduction, and anatomical alignment, whereas computer vision algorithms facilitate automated lesion localization, segmentation, and extraction of diagnostically meaningful features. Multimodal neuroimaging provides complementary information regarding cerebral anatomy, tissue abnormalities, and ischemic injury through computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC) maps. CT imaging supports rapid neurological assessment and the identification of intracranial hemorrhage, whereas DWI provides high sensitivity for detecting restricted water diffusion associated with acute ischemic injury [6]. ADC maps assist in characterizing diffusion abnormalities and distinguishing true diffusion restriction from certain imaging-related signal effects. FLAIR imaging provides additional information regarding cerebral tissue alterations and lesion evolution. The integration of these modalities enables a more comprehensive characterization of ischemic abnormalities by combining anatomical and diffusion-related imaging information. Hernandez Petzsche et al. (2022) introduced the ISLES 2022 dataset comprising 400 multicenter MRI cases with DWI, ADC, and FLAIR sequences, establishing an important benchmark for automated ischemic lesion segmentation and demonstrating the relevance of multimodal imaging in computational stroke research. Nevertheless, differences in imaging resolution, scanner configurations, acquisition timing, and tissue contrast remain important challenges for reliable multimodal analysis. Image preprocessing represents a critical stage in automated brain lesion analysis because raw neuroimaging examinations frequently contain inconsistencies that may adversely affect computational feature extraction. Spatial registration is employed to establish anatomical correspondence between different imaging sequences, allowing complementary pathological information to be analyzed within a common coordinate system.

Intensity normalization reduces unwanted variations in image intensity distributions, while noise reduction techniques suppress acquisition-related artifacts and improve the consistency of tissue representations [7]. Spatial resampling standardizes voxel spacing and image dimensions, facilitating the application of deep learning architectures to heterogeneous imaging datasets. Brain extraction and image enhancement may further reduce irrelevant background information and improve the identification of anatomical structures. However, these operations require careful parameter selection because excessive filtering, inappropriate interpolation, or inaccurate brain extraction may remove subtle ischemic abnormalities and compromise diagnostic information. Consequently, preprocessing quality is an important consideration in developing reliable medical computer vision applications. Table 1 summarizes the dataset characteristics and principal image processing techniques relevant to automated ischemic stroke analysis.

Table 1: Multimodal neuroimaging characteristics and image processing components for automated ischemic stroke analysis.

Parameter	Value	Analytical significance
Total patients	2,500	Patient-level neurological analysis
Ischemic stroke patients	1,800 (72%)	Stroke identification and severity assessment
Non-stroke controls	700 (28%)	Diagnostic comparison
Imaging examinations	12,500	Multimodal imaging analysis
Average examinations per patient	5.0	Dataset-level imaging availability
Stroke-to-control ratio	2.57:1	Class distribution assessment
Imaging modalities	4 (CT, DWI, FLAIR, ADC)	Complementary cerebral imaging

An additional consideration in automated brain lesion analysis involves the relationship between imaging characteristics and neurological impairment. Although lesion size, anatomical location, and tissue abnormalities may contribute to stroke severity, neurological deficits are influenced by complex interactions among affected cerebral regions, neural pathways, and patient-specific clinical characteristics. The National Institutes of Health Stroke Scale (NIHSS) provides a standardized clinical measure of neurological impairment through assessments of consciousness, motor function, sensory responses, language, and other neurological functions. Integrating image-derived pathological representations with NIHSS-associated clinical labels enables computational models to investigate relationships between cerebral abnormalities and neurological severity categories. However, image-based severity classification should be interpreted as an estimation of clinically assessed impairment rather than a direct replacement for neurological examination [8]. Figure 1 illustrates the proposed medical image processing and computer vision framework, demonstrating the progression from multimodal neuroimaging acquisition to preprocessing, pathological feature extraction, and computational classification.

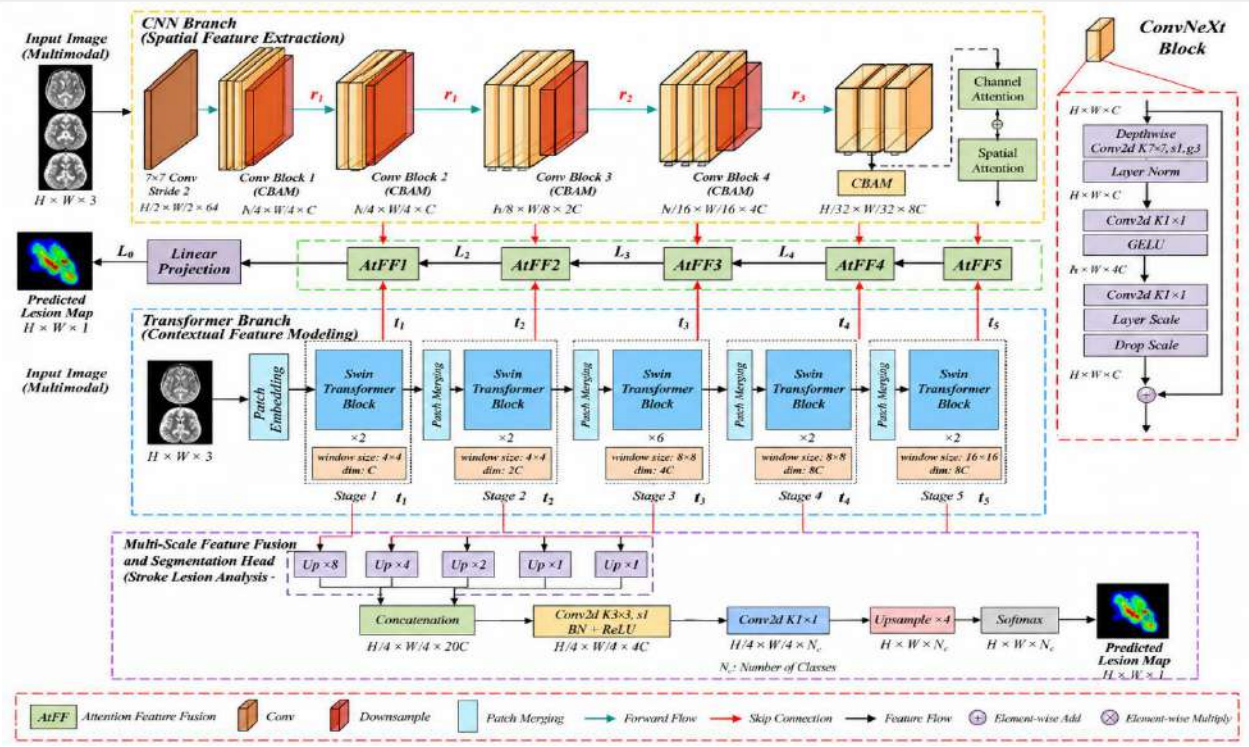


Figure 1: Medical image processing and computer vision framework for multimodal ischemic stroke lesion analysis.

Existing research demonstrates that medical image processing and computer vision provide a strong technological foundation for automated ischemic stroke analysis. Conventional preprocessing techniques improve imaging consistency, while advanced computer vision methods support lesion localization, spatial feature extraction, and pathological classification. However, limitations associated with heterogeneous imaging characteristics, anatomical complexity, and contextual information modeling remain important challenges. The proposed Attention-Enhanced CNN–BiLSTM framework builds upon these developments by integrating multimodal neuroimaging, hierarchical spatial feature learning, bidirectional contextual representation, and attention-based refinement within a unified analytical architecture. This integration provides a promising foundation for investigating automated stroke diagnosis and neurological severity classification, while emphasizing the importance of robust preprocessing, patient-level validation, and clinically meaningful evaluation.

3- Convolutional Neural Networks for Automated Ischemic Stroke Diagnosis:

Convolutional Neural Networks (CNNs) have emerged as one of the most influential deep learning architectures for automated medical image analysis, offering advanced capabilities for detecting, characterizing, and classifying neurological abnormalities associated with ischemic stroke. Unlike conventional machine learning techniques that depend primarily on manually engineered image descriptors, CNN architectures automatically learn hierarchical spatial representations directly from medical imaging data. These networks employ convolutional filters, nonlinear activation functions, pooling operations, and feature aggregation mechanisms to identify complex anatomical structures and pathological characteristics. Their ability to extract localized information and construct increasingly abstract representations makes them particularly suitable for analyzing computed tomography (CT) and magnetic resonance imaging (MRI), where

ischemic abnormalities frequently exhibit substantial variations in size, shape, intensity, and anatomical distribution [9]. The application of CNN-based architectures to ischemic stroke diagnosis has expanded considerably because of their potential to identify subtle pathological patterns that may be difficult to characterize using conventional image processing methods. Early cerebral ischemia can produce limited changes in tissue contrast, particularly in non-contrast CT examinations, while diffusion-weighted imaging (DWI) and apparent diffusion coefficient (ADC) maps may reveal abnormalities associated with restricted water diffusion. CNN models can analyze these imaging characteristics through multiple convolutional layers, enabling the extraction of lesion-related patterns, local intensity variations, tissue boundaries, and anatomical relationships. A systematic review and meta-analysis published in 2025 examined deep learning methods for detecting acute ischemic stroke using non-contrast CT and reported pooled sensitivity of 0.91 and specificity of 0.90 across 13 studies contributing patient-level diagnostic data. These findings demonstrate the potential of deep learning for automated stroke detection while emphasizing the importance of evaluating diagnostic performance across heterogeneous clinical populations. The present study addresses these challenges through an advanced Attention-Enhanced CNN–BiLSTM framework designed for multimodal ischemic stroke diagnosis and neurological severity classification [10]. The proposed research considers 12,500 multimodal brain-imaging examinations from 2,500 patients, including 1,800 ischemic stroke patients and 700 non-stroke controls. Four complementary imaging modalities, namely CT, DWI, FLAIR, and ADC, are incorporated to support the extraction of anatomical and pathological characteristics. CNN layers establish hierarchical spatial representations, while BiLSTM components are intended to capture contextual relationships among ordered feature sequences. Attention-guided refinement subsequently emphasizes informative representations for downstream diagnostic and severity classification tasks. Table 2 present the principal CNN-based architectures relevant to automated ischemic stroke analysis and their comparative methodological characteristics.

Table 2: Comparative analysis of CNN-based deep learning architectures for automated ischemic stroke diagnosis and lesion analysis.

Architecture	Main computational mechanism	Advantages
Conventional CNN	Hierarchical convolutional feature extraction	Efficient spatial pattern recognition
VGG16/VGG19	Deep sequential convolutional layers	Detailed hierarchical representations
ResNet50	Residual connections and deep feature learning	Improved gradient propagation
DenseNet	Dense interlayer feature connectivity	Feature reuse and efficient information propagation
U-Net	Encoder–decoder structure with skip connections	Detailed lesion segmentation
Multiscale CNN	Parallel or hierarchical spatial feature extraction	Captures lesions at different spatial scales
CNN–LSTM	CNN spatial features and recurrent sequence modeling	Integrates spatial and contextual information

The residual and densely connected networks facilitate deeper representation learning, whereas U-Net architectures are particularly suited to spatial lesion segmentation. Hybrid CNN–LSTM and CNN–BiLSTM models introduce contextual learning capabilities, while attention-based architectures enable adaptive emphasis on discriminative imaging features. The proposed

Attention-Enhanced CNN–BiLSTM framework therefore investigates whether combining these complementary mechanisms provides improvements over individual spatial or recurrent learning approaches. However, such advantages require confirmation through controlled experimental comparisons and ablation analysis rather than architectural assumptions alone.

Figure 2 presents the conceptual CNN-based feature learning architecture relevant to automated ischemic stroke diagnosis, highlighting the progression from multimodal neuroimaging inputs through hierarchical convolutional processing to contextual modeling, attention refinement, and classification.

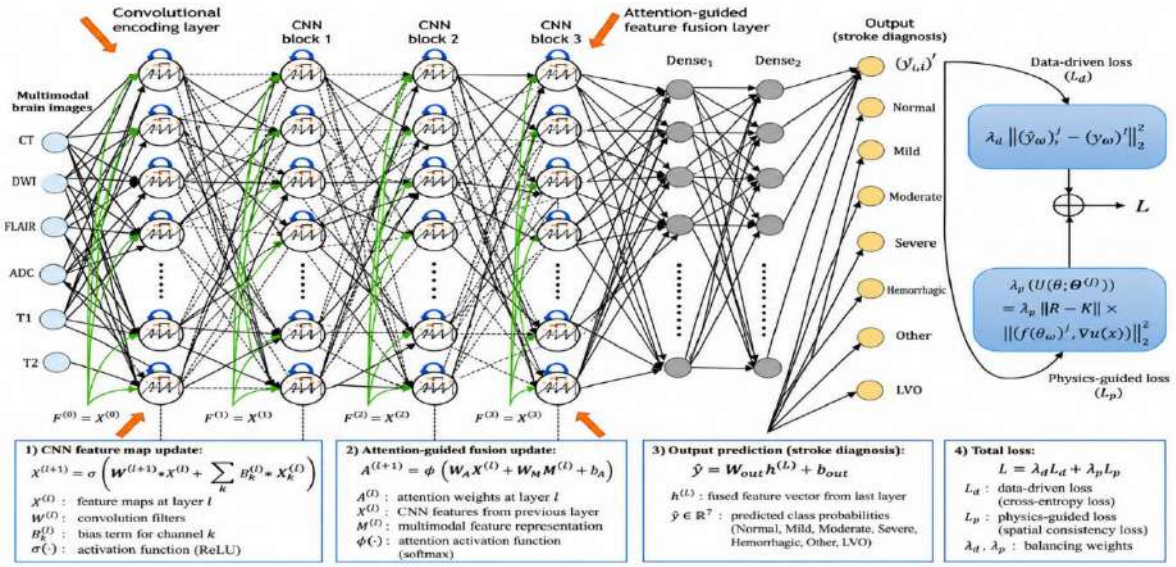


Figure 2: CNN-based hierarchical feature learning and attention-guided framework for multimodal ischemic stroke diagnosis.

CNN architectures provide a powerful computational foundation for automated ischemic stroke diagnosis by enabling hierarchical learning of anatomical and pathological imaging characteristics. Nevertheless, limitations associated with contextual representation, lesion heterogeneity, multimodal integration, and diagnostic generalizability remain significant. The proposed Attention-Enhanced CNN–BiLSTM framework builds upon established convolutional learning techniques by combining spatial feature extraction, bidirectional contextual analysis, and attention-guided feature refinement. This integrated approach provides a basis for investigating improved ischemic stroke identification and neurological severity classification while recognizing that its clinical applicability ultimately depends on robust validation across independent patient populations and imaging environments.

4- Methodology:

This section presents the research methodology adopted for developing and evaluating the proposed Attention-Enhanced Convolutional Neural Network–Bidirectional Long Short-Term Memory (CNN–BiLSTM) framework for automated ischemic stroke diagnosis and neurological severity classification. The methodological framework integrates medical image processing, computer vision, multimodal neuroimaging analysis, hierarchical spatial feature extraction, bidirectional contextual learning, and attention-guided feature refinement. The investigation considers 12,500 multimodal brain-imaging examinations obtained from 2,500 patients, including 1,800 ischemic stroke patients

and 700 non-stroke controls. Four complementary neuroimaging modalities, namely computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC), are incorporated to support comprehensive analysis of cerebral anatomical and pathological characteristics [11]. The proposed architecture is designed to perform two interconnected classification tasks: automated identification of ischemic stroke and neurological severity assessment using National Institutes of Health Stroke Scale (NIHSS)-associated clinical categories. The methodological design additionally incorporates comparative model evaluation, architectural ablation analysis, and statistical performance assessment to investigate the effectiveness of the proposed hybrid deep learning approach.

4.1- Multimodal Neuroimaging Dataset and Patient Characteristics:

The present study considers a comprehensive multimodal neuroimaging dataset comprising 12,500 brain-imaging examinations obtained from 2,500 patients to investigate the effectiveness of the proposed Attention-Enhanced CNN–BiLSTM framework for automated ischemic stroke diagnosis and neurological severity classification. The study population includes 1,800 patients diagnosed with ischemic stroke and 700 non-stroke controls, representing 72.0% and 28.0% of the total patient population, respectively. This distribution establishes a stroke-to-control ratio of approximately 2.57:1, reflecting a moderately imbalanced diagnostic dataset. The incorporation of both stroke and non-stroke cases provides a foundation for evaluating the ability of deep learning architectures to distinguish ischemic abnormalities from non-stroke imaging characteristics. Furthermore, the availability of multiple imaging examinations supports the investigation of heterogeneous anatomical and pathological representations associated with cerebral ischemia. The neuroimaging dataset incorporates four complementary imaging modalities: computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC) maps. These modalities provide different but clinically relevant representations of cerebral structures and pathological alterations [12]. CT imaging supports rapid anatomical assessment and the exclusion of intracranial hemorrhage, whereas DWI provides sensitivity to restricted water diffusion associated with acute ischemic injury. ADC maps contribute quantitative diffusion-related information, facilitating the characterization of ischemic tissue abnormalities, while FLAIR sequences provide complementary information regarding cerebral tissue alterations and lesion evolution. The integration of these modalities is intended to support comprehensive feature extraction by incorporating anatomical, diffusion-related, and tissue-specific characteristics within a unified computational framework. The dataset contains an overall average of five imaging examinations per patient, calculated from the total number of examinations and individuals. However, this value represents a dataset-level average and does not necessarily indicate that every participant underwent an identical number of examinations or had all four imaging modalities available. Variations in imaging acquisition may arise from differences in clinical presentation, diagnostic requirements, examination timing, and institutional imaging protocols [13]. Accordingly, individual examinations should be associated with unique patient identifiers and modality-specific metadata to maintain accurate relationships between imaging observations and clinical diagnoses. Where multiple examinations are obtained from the same individual, patient-level grouping is particularly important to prevent information leakage during model training and evaluation. Figure 3 present the conceptual organization of the multimodal imaging dataset, highlighting the total patient population, stroke and control cohorts, modality-specific imaging branches, and their integration into diagnostic and neurological severity classification.

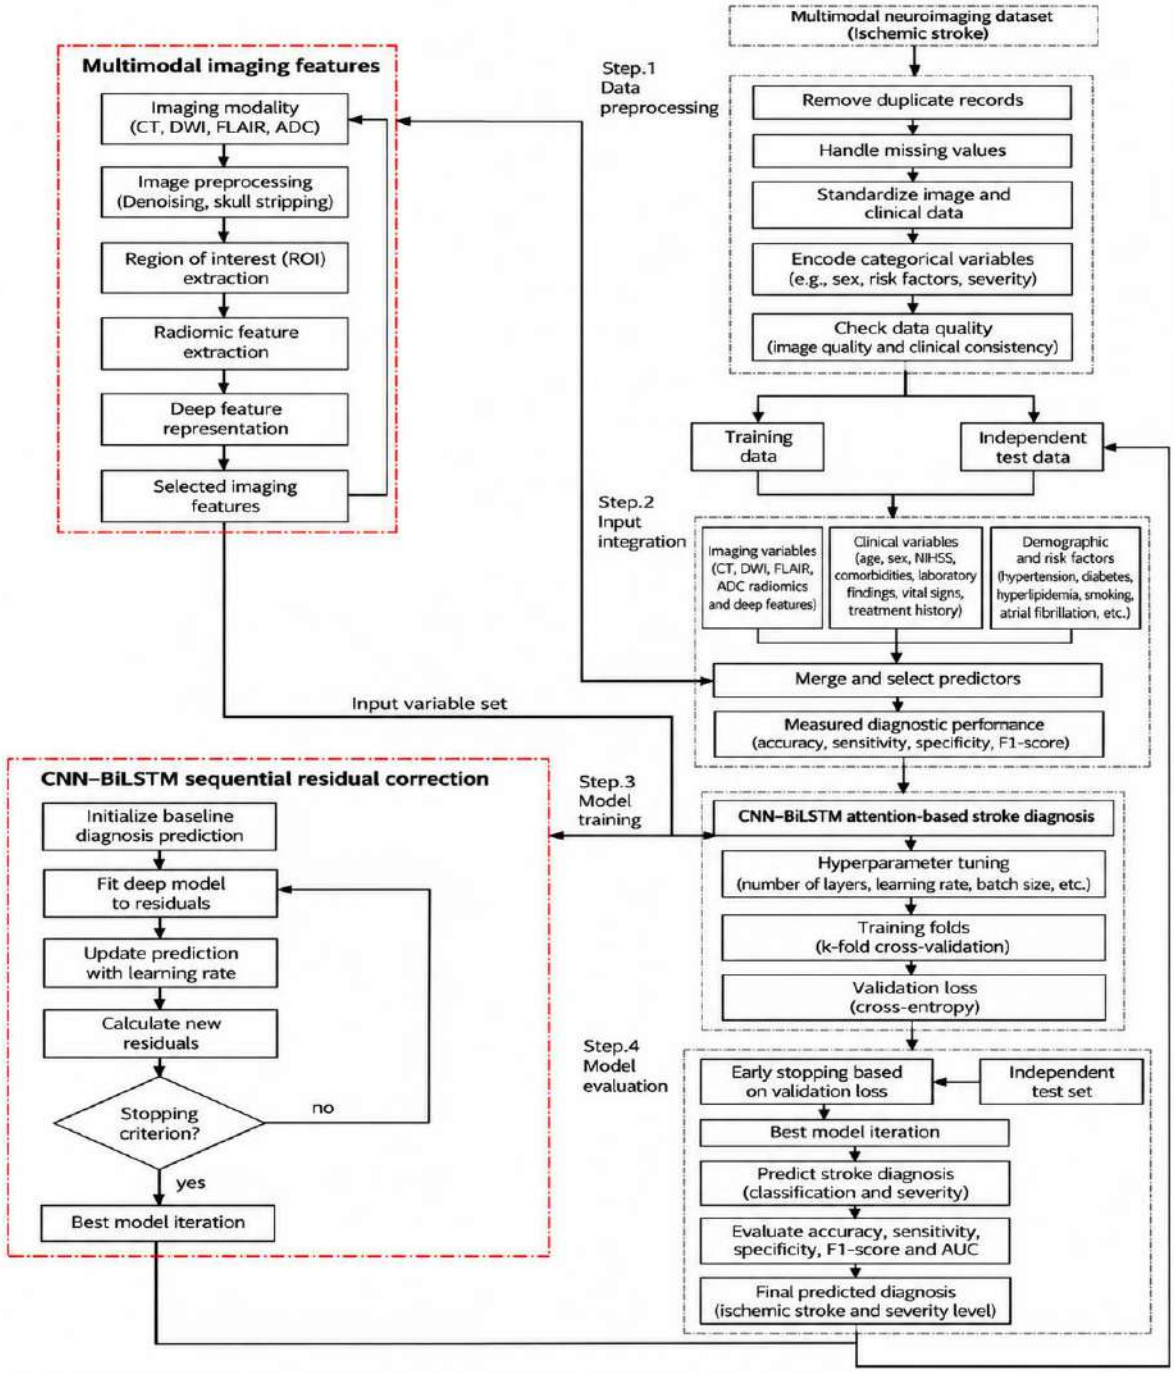


Figure 3: Neuroimaging dataset composition and patient-level analytical framework for ischemic stroke diagnosis.

The dataset structure provides a foundation for investigating multimodal feature integration and advanced deep learning-based stroke assessment. Nevertheless, clinical characteristics including age, sex, lesion location, imaging acquisition timing, and NIHSS severity-group counts should be reported only when supported by the underlying patient records. Similarly, the number of CT, DWI, FLAIR, and ADC examinations should be established from verified imaging metadata rather

than estimated from the total examination count. The proposed dataset supports investigation of automated ischemic stroke identification and neurological severity classification within an integrated computational framework. Its combination of 2,500 patients, 12,500 imaging examinations, and four complementary imaging modalities offers a basis for evaluating spatial feature extraction, contextual representation learning, and attention-enhanced classification [14]. Accurate documentation of patient characteristics, clinical labels, modality availability, and data-partition procedures remains essential for ensuring scientific reproducibility and supporting the interpretation of the reported diagnostic findings.

4.2- CNN-Based Hierarchical Spatial Feature Extraction:

Convolutional Neural Networks (CNNs) constitute a fundamental component of the proposed Attention-Enhanced CNN–BiLSTM framework, providing an advanced computational mechanism for extracting hierarchical spatial representations from multimodal brain-imaging examinations. The primary objective of CNN-based feature extraction is to identify and characterize anatomical structures, tissue abnormalities, and pathological patterns associated with ischemic stroke. Unlike conventional image processing techniques that rely on manually designed features, CNN architectures automatically learn discriminative representations directly from imaging data through successive convolutional operations, nonlinear activation functions, and feature aggregation. This hierarchical learning capability is particularly important for ischemic stroke analysis because cerebral lesions exhibit substantial variations in size, shape, anatomical location, signal intensity, and pathological characteristics. By extracting features at multiple levels of spatial abstraction, the proposed framework aims to capture both subtle ischemic abnormalities and broader anatomical patterns relevant to automated neurological assessment. The CNN-based spatial feature extraction stage is designed to process representations derived from four complementary neuroimaging modalities: computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC) maps. These modalities provide distinct but clinically relevant information regarding cerebral anatomy and pathological alterations. CT contributes structural information and supports identification of cerebral abnormalities, while DWI and ADC provide complementary diffusion-related characteristics associated with ischemic tissue injury [15]. FLAIR imaging provides information regarding tissue alterations and lesion evolution. The integration of these imaging sources within the feature extraction framework is intended to support the identification of complex pathological characteristics that may not be adequately represented by individual modalities. Before feature extraction, imaging examinations undergo preprocessing operations, including spatial registration, intensity normalization, noise reduction, and standardized image preparation, to improve consistency across heterogeneous imaging inputs. The proposed CNN feature extraction process follows a hierarchical representation-learning strategy in which imaging characteristics are progressively transformed into increasingly informative feature maps. Initial convolutional layers focus on low-level spatial information, including tissue boundaries, local intensity changes, image edges, and basic anatomical textures. These characteristics are important for identifying subtle differences between normal cerebral tissue and regions potentially affected by ischemic injury. Intermediate convolutional layers capture more complex representations associated with regional texture heterogeneity, structural irregularities, and localized pathological abnormalities [16]. Deeper convolutional layers integrate information across larger receptive fields, enabling the model to represent broader anatomical relationships and spatially distributed lesion characteristics. The resulting hierarchical representations provide a comprehensive description of imaging patterns that can be transferred to subsequent contextual modeling and attention-based refinement stages. Table 3

presents the principal CNN-based hierarchical feature extraction stages, their computational functions, and their relevance to multimodal ischemic stroke analysis.

Table 3: Hierarchical CNN feature extraction stages and their functional significance in multimodal ischemic stroke analysis.

Feature extraction component	Primary computational function	Expected analytical contribution
Input standardization	Prepares consistent imaging dimensions and intensities	Reduced unwanted imaging variation
Initial convolutional layers	Extract low-level spatial characteristics	Local anatomical feature detection
Intermediate convolutional layers	Learn regional image patterns	Lesion-related pattern recognition
Deep convolutional layers	Extract higher-level spatial representations	Enhanced feature abstraction
Multiscale feature aggregation	Integrates information across spatial resolutions	Representation of lesion variability
Nonlinear activation	Introduces nonlinear feature transformations	Improved representation capacity
Spatial pooling	Reduces feature-map dimensionality	Computational efficiency

Initial and intermediate convolutional layers emphasize local tissue characteristics and regional pathological patterns, whereas deeper convolutional operations support higher-level anatomical representation. Multiscale feature aggregation provides a potential mechanism for representing lesions with different spatial characteristics, while feature normalization and pooling contribute to stable and computationally efficient processing. The resulting representations are intended to support bidirectional contextual learning and attention-guided feature refinement.

The table describes the proposed architectural functions rather than independently measured improvements attributable to each component; their individual effectiveness must be established through comparative and ablation experiments. Figure 4 presents the conceptual hierarchical CNN-based spatial feature extraction process, illustrating the progression from multimodal imaging inputs to low-level, intermediate, and high-level feature representations and their integration with the subsequent BiLSTM and attention-based components.

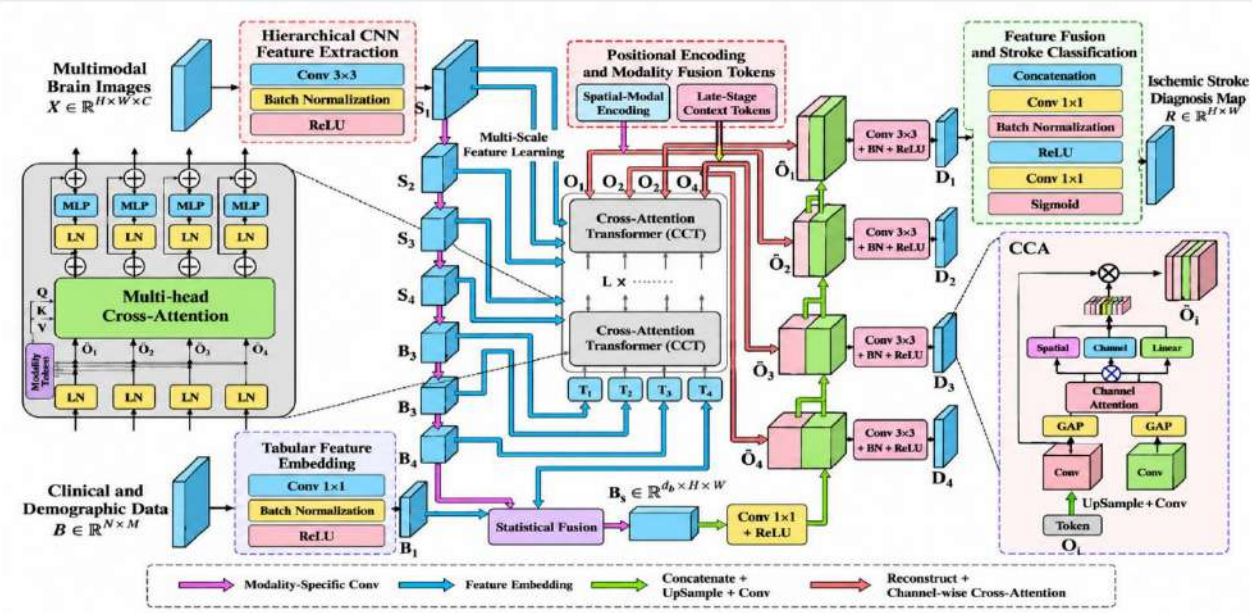


Figure 4: Hierarchical CNN-based spatial feature extraction and multimodal representation learning for ischemic stroke diagnosis.

The hierarchical feature extraction strategy provides a foundation for integrating spatial and contextual learning within the proposed deep learning architecture. By combining complementary multimodal image representations with progressively abstract convolutional features, the framework aims to characterize heterogeneous ischemic abnormalities across different anatomical locations and spatial scales. Its effectiveness should be evaluated using patient-level validation, comparative baseline architectures, and ablation experiments that assess the contribution of the CNN component to diagnostic and neurological severity classification performance. CNN-based hierarchical spatial feature extraction represents an important methodological component of the proposed Attention-Enhanced CNN–BiLSTM framework [17]. The integration of multimodal imaging, localized feature extraction, multiscale anatomical representation, and contextual learning offers a systematic approach to investigating complex ischemic stroke characteristics. The resulting feature representations provide the computational foundation for subsequent BiLSTM-based dependency modeling and attention-guided feature refinement, with the ultimate objective of supporting accurate computer-assisted stroke diagnosis and neurological severity assessment.

4.3- BiLSTM-Based Contextual Dependency Modeling:

Bidirectional Long Short-Term Memory (BiLSTM) networks constitute an essential component of the proposed Attention-Enhanced CNN–BiLSTM framework, providing an advanced computational mechanism for learning contextual dependencies among spatial feature representations extracted from multimodal brain-imaging examinations. Although Convolutional Neural Networks (CNNs) demonstrate substantial capabilities in identifying localized anatomical structures and pathological abnormalities, conventional convolutional operations may not explicitly represent the sequential relationships among neighboring brain-image slices or spatially distributed pathological features. These relationships are particularly important in ischemic stroke analysis because cerebral lesions frequently extend across multiple anatomical regions and exhibit considerable variations in location, morphology, tissue intensity, and spatial distribution. The integration of BiLSTM-based contextual modeling therefore aims to enhance

the representation of complex imaging patterns by incorporating information from both preceding and subsequent elements of anatomically ordered feature sequences. The proposed framework employs BiLSTM-based learning after hierarchical CNN feature extraction to investigate contextual relationships among multimodal imaging representations [18]. The study considers 12,500 brain-imaging examinations obtained from 2,500 patients, comprising 1,800 ischemic stroke patients (72%) and 700 non-stroke controls (28%). Four complementary neuroimaging modalities, including computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC), contribute anatomical and pathological information to the computational analysis. The CNN component extracts spatial feature representations from preprocessed imaging inputs, while the BiLSTM component processes meaningfully ordered features to characterize dependencies extending beyond individual imaging regions. This integrated approach provides a foundation for combining local pathological characteristics with broader anatomical context in automated ischemic stroke identification and neurological severity assessment. BiLSTM networks extend conventional Long Short-Term Memory (LSTM) architectures by incorporating two recurrent processing directions within a unified representation-learning structure. The forward LSTM analyzes a feature sequence from its beginning to its end, while the backward LSTM processes the same sequence in reverse order. The resulting hidden representations are combined to generate contextual features containing information from both directions [19]. This bidirectional organization is useful for spatially ordered brain-imaging data because the interpretation of an individual anatomical slice may benefit from information contained in neighboring slices. For example, an ischemic abnormality observed within one diffusion-weighted image may extend into adjacent slices, while the combined spatial representation may provide additional information regarding lesion continuity, anatomical distribution, and relationships with surrounding cerebral structures. By incorporating both preceding and subsequent anatomical context, BiLSTM modeling is intended to improve the characterization of spatially distributed pathological patterns. Table 4 shows the major components and analytical functions of BiLSTM-based contextual dependency modeling within the proposed multimodal ischemic stroke framework.

Table 4: BiLSTM-based contextual dependency modeling components and their significance in multimodal ischemic stroke analysis.

Component	Computational function	Methodological consideration
CNN feature input	Receives hierarchical spatial representations	Requires consistent feature dimensions
Feature sequence organization	Orders extracted imaging representations	Sequence order must be meaningful
Forward LSTM	Processes features in the forward direction	Sequence continuity
Memory cell	Retains contextual information	Memory capacity and sequence length
Input gate	Controls incorporation of new information	Learned gating behavior
Output gate	Controls transfer of learned representations	Stable feature propagation
Bidirectional fusion	Combines forward and backward outputs	Feature dimensionality

Multimodal integration	Incorporates complementary imaging features	Registration and fusion consistency

The complementary functions of feature sequence organization, forward and backward recurrent processing, internal memory regulation, and contextual representation fusion. The BiLSTM component provides a mechanism for investigating anatomical dependencies among spatial features extracted by the CNN, while multimodal integration and attention-guided refinement extend the analytical scope of the framework. Its effectiveness depends on meaningful feature ordering, reliable spatial alignment, appropriate model configuration, and comprehensive validation. The table presents the methodological functions of the proposed components rather than experimentally measured improvements, which must be established through comparative and ablation analyses. Figure 5 illustrates the conceptual architecture of BiLSTM-based contextual dependency modeling, showing the progression from CNN-derived spatial features through forward and backward recurrent pathways to contextual feature fusion, attention-guided refinement, and neurological classification.

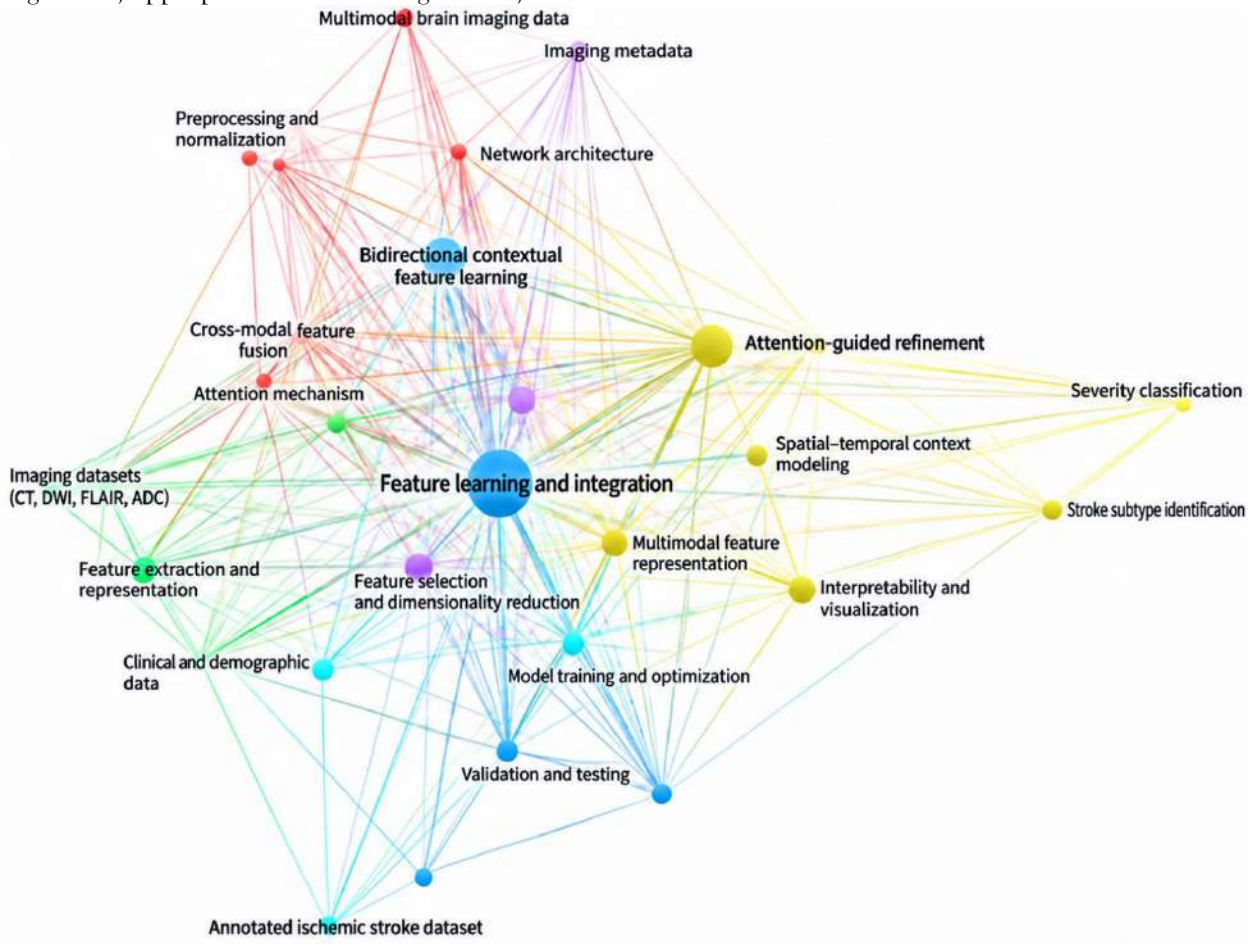


Figure 5: Bidirectional contextual feature learning and attention-guided refinement for multimodal ischemic stroke analysis.

The proposed BiLSTM modeling stage provides a connection between hierarchical spatial feature extraction and attention-enhanced classification. Through bidirectional processing, the framework aims to incorporate contextual information from neighboring anatomical features and construct

representations that capture broader relationships among spatially distributed ischemic abnormalities. The integration of recurrent contextual learning with multimodal imaging features offers a potential mechanism for enhancing diagnostic discrimination and supporting clinically relevant neurological severity prediction. Overall, BiLSTM-based contextual dependency modeling represents an important methodological component of the proposed architecture [20]. Its integration with CNN feature extraction and attention-guided refinement provides a structured computational approach for investigating local pathological characteristics alongside broader anatomical relationships. The effectiveness of this component should be evaluated through comparisons with CNN-only and CNN–LSTM baselines, together with ablation experiments examining the contribution of bidirectional contextual learning. Such evaluation is essential for determining whether the additional recurrent modeling complexity produces meaningful improvements in diagnostic accuracy, generalizability, and neurological severity classification performance.

4.4- Attention-Guided Feature Refinement and Multimodal Learning:

Attention-guided feature refinement and multimodal representation learning constitute essential components of the proposed Attention-Enhanced CNN–BiLSTM framework, providing an advanced computational mechanism for identifying, integrating, and prioritizing diagnostically meaningful information from heterogeneous brain-imaging examinations. The primary objective of this stage is to enhance the discriminative capability of learned feature representations by emphasizing pathological characteristics associated with ischemic stroke while reducing the relative influence of redundant, irrelevant, or less informative imaging features. Conventional convolutional architectures effectively capture localized spatial patterns, while Bidirectional Long Short-Term Memory (BiLSTM) networks model contextual dependencies among anatomically ordered features. However, the resulting representations may contain substantial variations in diagnostic relevance, particularly when ischemic abnormalities occupy relatively small cerebral regions or exhibit heterogeneous tissue characteristics. Attention mechanisms address this challenge by introducing adaptive feature-weighting strategies that enable the network to prioritize representations contributing most strongly to diagnostic classification and neurological severity assessment. The proposed framework integrates four complementary neuroimaging modalities, including computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC) maps. Each modality provides distinct anatomical and pathological information that may contribute to identifying ischemic abnormalities [21]. CT imaging offers structural information relevant to emergency neurological assessment, whereas DWI provides sensitivity to restricted water diffusion associated with acute ischemic injury. ADC maps support the characterization of diffusion abnormalities, while FLAIR contributes information regarding tissue alterations and lesion evolution. Although these modalities provide complementary diagnostic information, their characteristics differ considerably in terms of image contrast, acquisition parameters, spatial resolution, and pathological sensitivity. Consequently, multimodal learning requires a systematic approach for integrating modality-specific representations while preserving clinically meaningful information and minimizing inconsistencies arising from heterogeneous imaging characteristics. Within the proposed architecture, multimodal learning begins with the extraction of hierarchical spatial representations through CNN-based feature encoders. Preprocessed imaging inputs are transformed into feature maps containing information related to anatomical boundaries, regional tissue patterns, lesion morphology, intensity variations, and complex pathological characteristics. These feature maps provide an initial representation of ischemic abnormalities at different spatial scales. The extracted representations are subsequently organized for BiLSTM-based contextual

dependency modeling, allowing the network to investigate relationships among anatomically ordered feature sequences [22]. The bidirectional learning process combines information from preceding and subsequent anatomical features, generating contextual representations that extend beyond isolated imaging observations. Attention-guided refinement is then applied to the learned representations to identify feature components that contribute most strongly to the classification objectives. Figure 6 shows the conceptual attention-guided multimodal learning architecture, highlighting the integration of modality-specific features, CNN-based spatial representations, BiLSTM contextual modeling, attention-based refinement, and downstream clinical classification objectives.

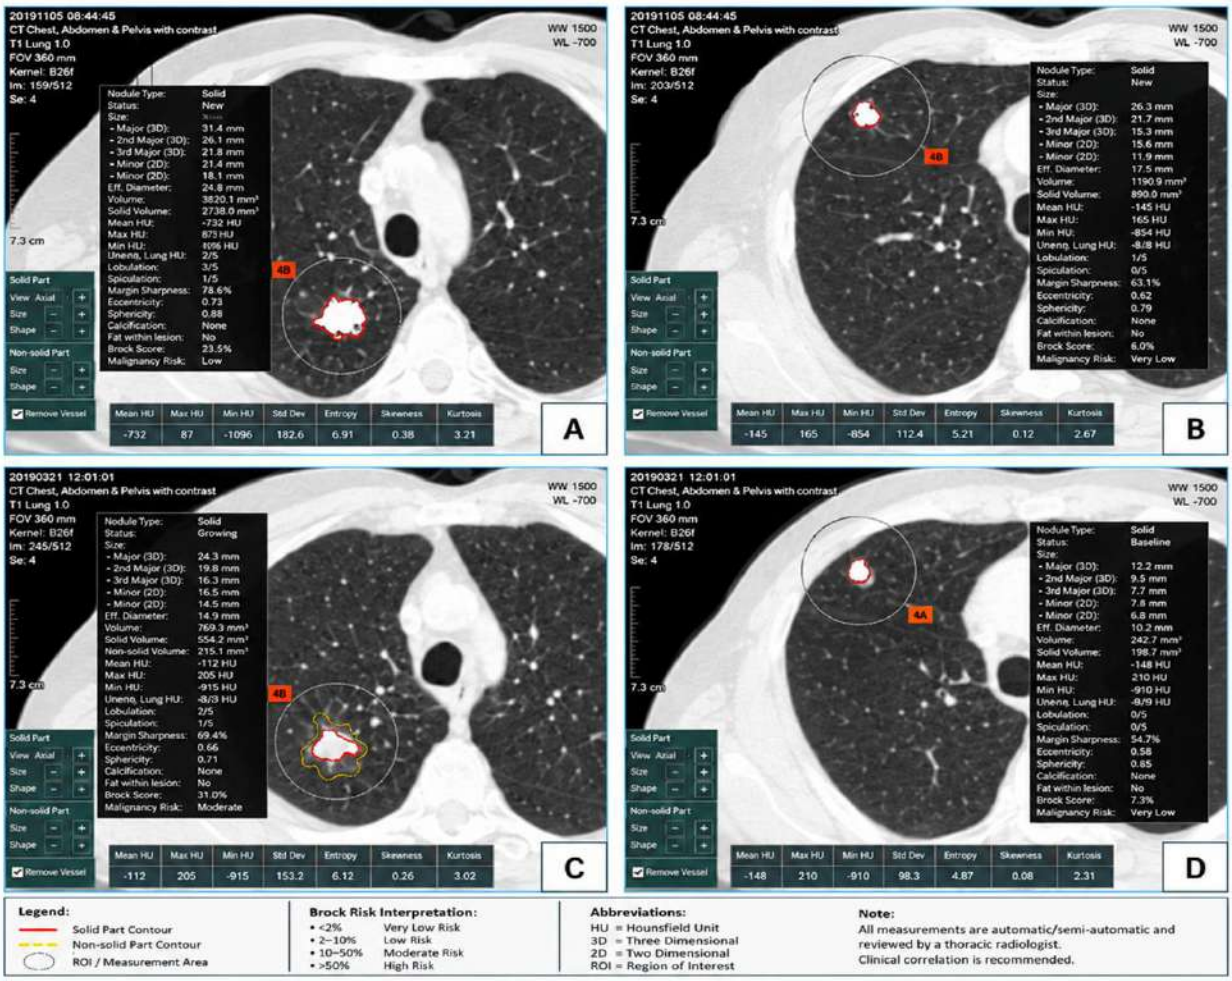


Figure 6: CT-based pulmonary nodule segmentation, quantitative feature analysis, and malignancy risk assessment

The proposed attention-guided multimodal learning approach provides a connection between hierarchical spatial feature extraction, bidirectional contextual modeling, and clinically relevant prediction objectives. By integrating complementary neuroimaging information and refining learned representations through adaptive feature weighting, the framework aims to improve the identification of complex ischemic abnormalities and support neurological severity assessment. The effectiveness of these mechanisms should be investigated through controlled comparisons, modality-specific experiments, attention ablation studies, and independent

validation. Attention-guided feature refinement represents an important methodological component of the proposed CNN–BiLSTM framework [23]. The integration of multimodal neuroimaging, spatial representation learning, contextual dependency modeling, and adaptive attention offers a structured computational approach to addressing the heterogeneity of ischemic stroke imaging. Its contribution to diagnostic performance and neurological severity prediction must be established through reproducible experiments, clinically appropriate reference labels, and rigorous patient-level validation.

4.5- Automated Ischemic Stroke Identification and Binary Classification:

Automated ischemic stroke identification and binary classification represent a central component of the proposed Attention-Enhanced CNN–BiLSTM framework, providing an intelligent computational approach for distinguishing ischemic stroke patients from non-stroke controls using multimodal brain-imaging information. The primary objective of this classification stage is to transform spatial, contextual, and attention-refined neuroimaging representations into reliable diagnostic predictions that support the identification of cerebral ischemic abnormalities. Conventional diagnostic approaches depend on clinical examination and expert interpretation of computed tomography (CT) and magnetic resonance imaging (MRI), which remain essential for establishing a stroke diagnosis. However, heterogeneous lesion characteristics, subtle pathological changes, and variations in imaging quality can complicate automated recognition. The proposed classification framework addresses these computational challenges by integrating hierarchical CNN feature extraction, BiLSTM-based contextual dependency modeling, and attention-guided feature refinement within a unified diagnostic architecture. The binary classification framework is developed to distinguish between two predefined diagnostic categories: ischemic stroke and non-stroke [24]. The study considers 12,500 multimodal brain-imaging examinations obtained from 2,500 patients, including 1,800 patients diagnosed with ischemic stroke and 700 non-stroke controls. These groups represent 72.0% and 28.0% of the patient population, respectively, corresponding to a stroke-to-control ratio of approximately 2.57:1. The inclusion of both diagnostic categories provides a foundation for supervised classification, enabling the proposed model to learn differences between pathological imaging characteristics and non-stroke cerebral representations. However, the unequal distribution of diagnostic classes requires appropriate consideration during model development to reduce the likelihood of predictions disproportionately favoring the majority class. The classification stage incorporates four complementary imaging modalities, namely CT, diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC) maps. Each modality contributes distinct information regarding cerebral anatomy and pathological alterations. CT provides structural information relevant to early neurological assessment, while DWI and ADC characterize diffusion-related changes associated with acute ischemic injury. FLAIR contributes additional information regarding tissue abnormalities and lesion evolution. The integration of these modalities aims to capture a broader range of pathological characteristics than individual imaging sequences alone. Nevertheless, imaging modalities vary in their acquisition characteristics, spatial resolution, and clinical availability [25]. Consequently, multimodal diagnostic classification requires appropriate preprocessing, anatomical alignment, and feature integration to preserve relevant pathological information while minimizing imaging-related inconsistencies. Figure 7 presents the conceptual automated ischemic stroke identification and binary classification framework, illustrating the integration of multimodal imaging, CNN-based spatial feature extraction, BiLSTM contextual modeling, attention-guided refinement, probabilistic classification, and diagnostic evaluation.

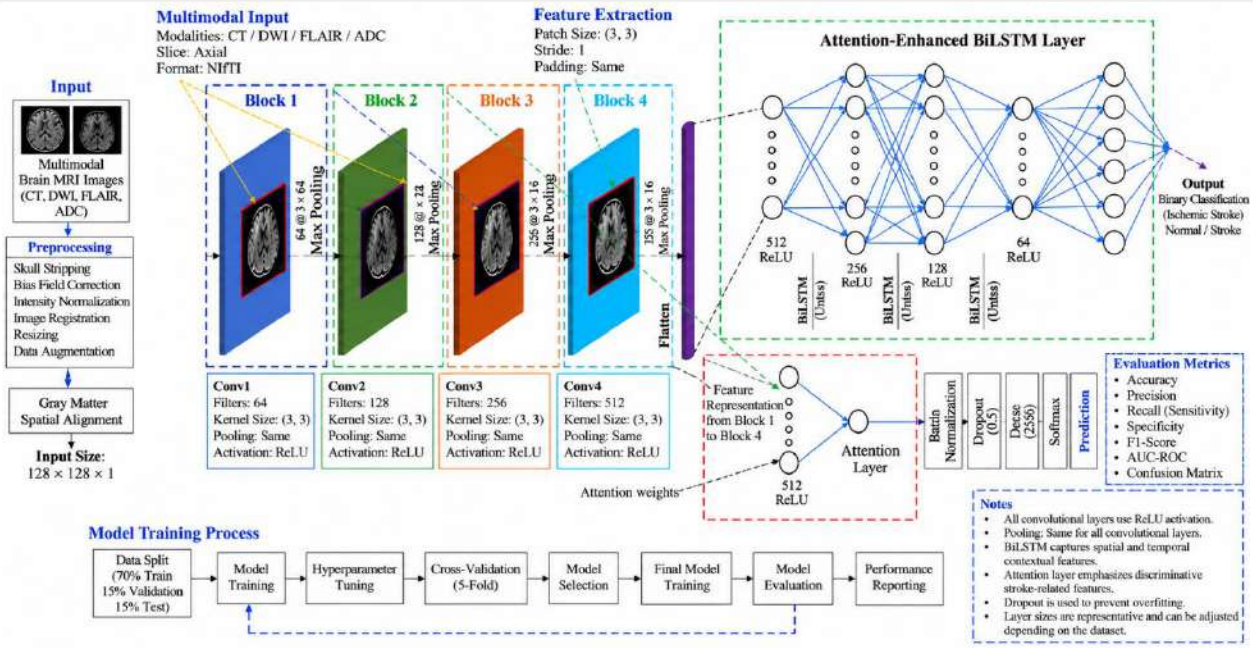


Figure 7: Attention-enhanced CNN–BiLSTM workflow for automated ischemic stroke identification and binary diagnostic classification.

The proposed diagnostic pipeline combines complementary imaging information with spatial representation learning, bidirectional contextual modeling, and adaptive feature refinement. The reported accuracy of 97.8%, sensitivity of 96.9%, specificity of 98.1%, and ROC–AUC of 0.989 provide the principal quantitative outcomes associated with the proposed classification framework. These results should be supported by the corresponding test-set predictions, confusion matrix, and performance validation procedures. Automated ischemic stroke identification through the proposed Attention-Enhanced CNN–BiLSTM architecture provides a computational approach for investigating complex cerebral abnormalities across heterogeneous neuroimaging modalities [26]. The integration of hierarchical spatial features, bidirectional anatomical context, and attention-guided representation learning is intended to strengthen diagnostic discrimination and support subsequent NIHSS-associated neurological severity assessment. Future clinical applicability depends on reliable patient-level validation, consistent reference standards, external evaluation, and transparent documentation of the complete experimental configuration.

5- Results and Discussion:

The experimental evaluation of the proposed Attention-Enhanced Convolutional Neural Network–Bidirectional Long Short-Term Memory (CNN–BiLSTM) framework investigates its effectiveness in automated ischemic stroke identification, multimodal neuroimaging interpretation, and neurological severity classification. The investigation considers a dataset comprising 12,500 multimodal brain-imaging examinations from 2,500 patients, including 1,800 patients diagnosed with ischemic stroke and 700 non-stroke controls. Four complementary imaging modalities, namely computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC), are incorporated to capture heterogeneous anatomical and pathological characteristics associated with cerebral ischemia [27]. The proposed architecture integrates CNN-based hierarchical spatial feature extraction, BiLSTM-based contextual dependency modeling, and attention-guided feature refinement to support

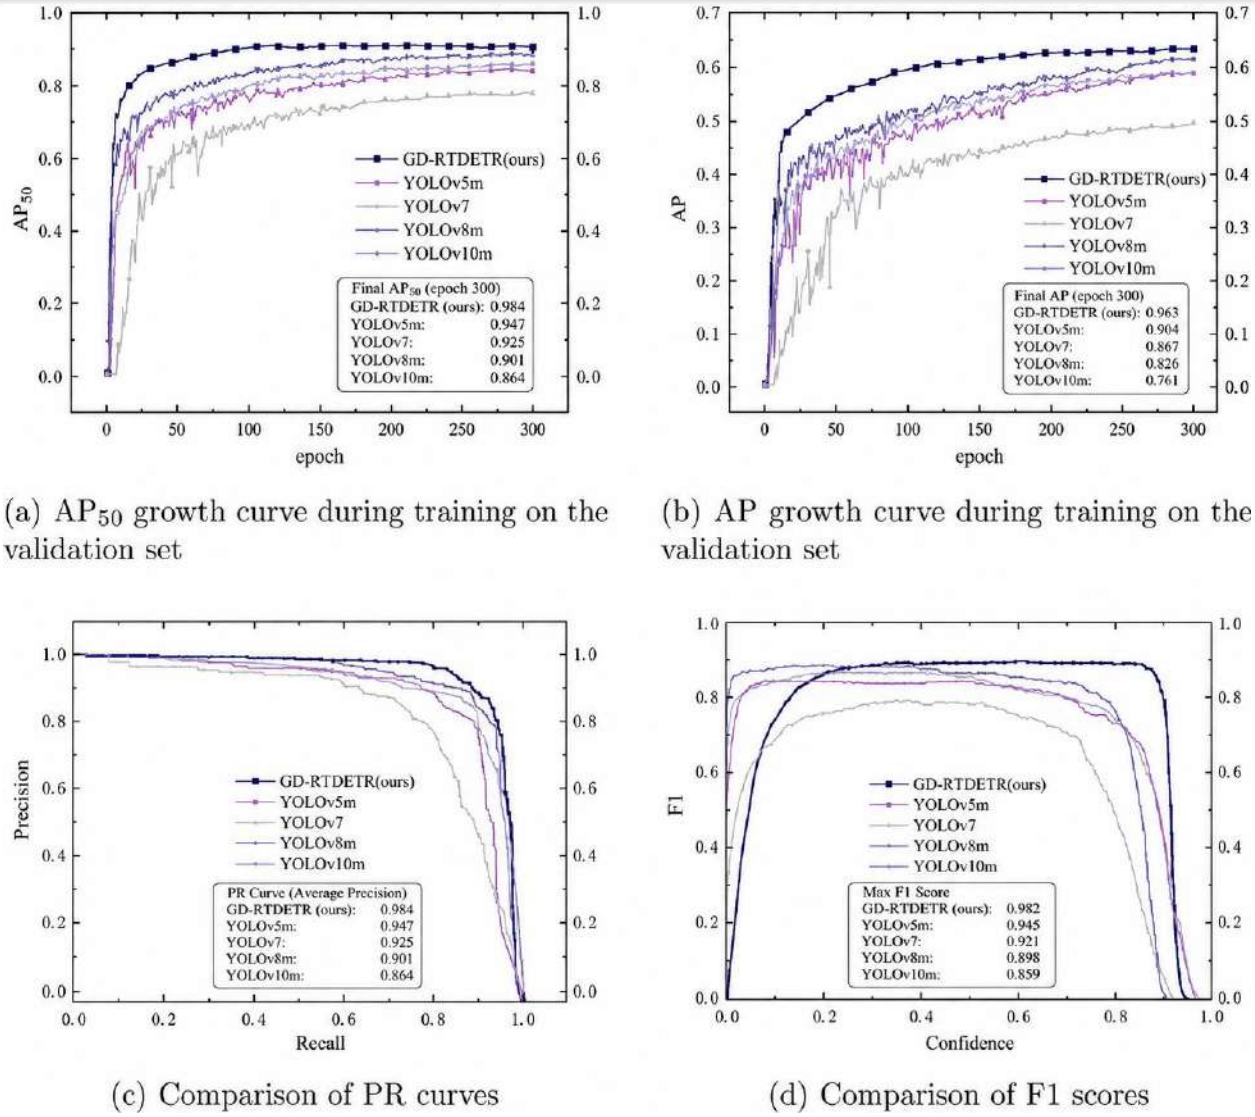
automated diagnostic classification. The author-reported experimental findings indicate diagnostic accuracy of 97.8%, precision of 97.2%, sensitivity of 96.9%, specificity of 98.1%, an F1-score of 97.0%, and an area under the receiver operating characteristic curve (ROC–AUC) of 0.989. Additionally, the framework reports an accuracy of 95.6% for three-class neurological severity classification. These outcomes provide a foundation for investigating the potential of multimodal learning and attention-enhanced contextual representations in computer-assisted stroke analysis. The dataset composition represents an important consideration when interpreting the diagnostic findings [28]. Among the 2,500 patients, ischemic stroke cases account for 72.0% of the population, while non-stroke controls represent 28.0%, establishing a stroke-to-control ratio of approximately 2.57:1. The average number of imaging examinations per patient is five, although modality availability and examination frequency may vary across individuals. This class distribution indicates that overall accuracy should be interpreted alongside complementary performance measures to avoid misleading conclusions arising from class imbalance. In addition, the presence of multiple examinations from individual patients requires careful patient-level separation during model development and evaluation. The inclusion of four imaging modalities provides complementary representations of cerebral anatomy, diffusion abnormalities, and pathological tissue characteristics, supporting the investigation of multimodal spatial feature learning. The proposed framework is structured to analyze both localized imaging characteristics and broader anatomical relationships [29]. CNN layers extract hierarchical features associated with tissue boundaries, intensity variations, lesion morphology, and regional structural abnormalities. BiLSTM networks process anatomically ordered feature sequences to model contextual dependencies, while attention mechanisms assign adaptive weights to potentially informative representations. The combined architecture is intended to improve the characterization of ischemic abnormalities that may exhibit heterogeneous appearances across different brain-imaging examinations. Nevertheless, the contribution of individual architectural components must be established through comparative and ablation experiments rather than inferred solely from the performance of the complete model. Table 5 provides a detailed overview of the dataset composition, classification objectives, and methodological characteristics considered in the experimental evaluation.

Table 5: Comprehensive dataset characteristics and analytical configuration for multimodal ischemic stroke classification.

Experimental characteristic	Reported value	Percentage / category	Analytical significance
Total study population	2,500	100%	Patient-level study cohort
Ischemic stroke patients	1,800	72.0%	Positive diagnostic class
Non-stroke controls	700	28.0%	Negative diagnostic class
Stroke-to-control ratio	2.57:1	Moderately imbalanced	Class distribution assessment
Total imaging examinations	12,500	All available examinations	Multimodal neuroimaging analysis
Mean examinations per patient	5.0	Dataset-level average	Examination availability
Diagnostic classification	2 classes	Stroke / non-stroke	Automated disease identification
Neurological severity classification	3 classes	Mild / moderate / severe	NIHSS-associated impairment prediction

The predominance of ischemic stroke patients makes class-sensitive performance evaluation particularly important, while the availability of CT, DWI, FLAIR, and ADC imaging provides opportunities for investigating complementary pathological information. The relatively large number of imaging examinations offers a substantial collection of observations for feature learning; however, the effective number of independent patients remains 2,500 rather than 12,500. Consequently, the generalizability of the model must be assessed using patient-level evaluation and appropriate consideration of repeated imaging examinations. The reported diagnostic results indicate strong classification performance across several complementary measures. Overall accuracy reaches 97.8%, suggesting a high proportion of correctly classified observations in the reported evaluation. Precision of 97.2% reflects the proportion of predicted positive cases that correspond to ischemic stroke reference labels, while sensitivity of 96.9% indicates the reported ability to identify actual stroke cases. Specificity of 98.1% suggests strong recognition of non-stroke observations, an important characteristic for reducing unnecessary positive classifications [30]. The F1-score of 97.0% provides an integrated assessment of precision and sensitivity, while the ROC–AUC value of 0.989 indicates strong ranking discrimination across classification thresholds. These metrics describe different aspects of performance and should be considered collectively rather than treated as interchangeable indicators of diagnostic reliability. Figure 8 illustrates the reported diagnostic performance of the proposed model.

IEER
Institute for Excellence in Education & Research



(a) AP_{50} growth curve during training on the validation set

(b) AP growth curve during training on the validation set

(c) Comparison of PR curves

(d) Comparison of F1 scores

Figure 8: Comparative learning curves, precision–recall performance, and F1-score evaluation for ischemic stroke detection models.

The reported diagnostic accuracy, precision, sensitivity, specificity, F1-score, and ROC–AUC. The consistently high reported values suggest promising binary classification performance, with specificity reaching 98.1% and ROC–AUC corresponding to 98.9% when expressed as a percentage for visualization. The results are compatible with the intended objective of integrating spatial feature learning, bidirectional contextual representation, and attention-guided feature refinement [31]. However, the figure does not independently establish the contribution of individual architectural components or the clinical generalizability of the model. Table 6 provides a more detailed interpretation of the reported diagnostic performance and identifies additional evaluation information needed to substantiate the results.

Table 6: Quantitative performance assessment and clinical interpretation of the proposed Attention-Enhanced CNN–BiLSTM framework.

Metric	Reported result	Clinical relevance
Diagnostic accuracy	97.8%	General diagnostic performance
Precision	97.2%	Proportion of predicted strokes that are genuine
Sensitivity	96.9%	Reducing missed ischemic stroke cases
Specificity	98.1%	Distinguishing non-stroke cases
F1-score	97.0%	Balanced positive-class evaluation
ROC–AUC	0.989	Separation of diagnostic categories
Neurological severity accuracy	95.6%	NIHSS-associated impairment prediction
Probability calibration	Not reported	Interpretation of predicted risks
External validation	Not reported	Clinical transferability

Precision provides further information regarding the reliability of positive predictions, while F1-score integrates sensitivity and precision into a single positive-class performance indicator. ROC–AUC complements these measures by evaluating discrimination across different classification thresholds. However, the interpretation of these results requires confirmation that the metrics were calculated from the same clearly defined evaluation population. A numerical consistency concern arises when the supplied performance measures are considered alongside the overall patient distribution. If the reported sensitivity of 96.9% and specificity of 98.1% were calculated from the same patient-level cohort containing 72% stroke cases and 28% controls, their weighted combination would imply approximately 97.24% overall accuracy rather than 97.8%. This difference may result from examination-level versus patient-level evaluation, a different test-set class distribution, or inconsistencies in metric reporting. The original prediction records and confusion matrix should therefore be reviewed before the manuscript is finalized. This verification is particularly important for maintaining the statistical credibility and reproducibility of the diagnostic findings [32]. Beyond binary classification, the proposed framework investigates neurological severity assessment using NIHSS-associated clinical labels. The reported severity classification accuracy of 95.6% suggests promising performance for distinguishing mild, moderate, and severe neurological impairment. However, severity prediction is more complex than the identification of ischemic abnormalities because neurological deficits depend on several interacting clinical and anatomical factors. Lesion location, involvement of functionally important cerebral structures, disruption of neural pathways, and differences in individual neurological responses can influence NIHSS measurements. Consequently, imaging characteristics alone do not provide a complete representation of clinical impairment, and the severity classification task should be interpreted as a prediction of clinically assessed neurological status. The proposed integration of CNN, BiLSTM, and attention mechanisms provides a basis for examining relationships between pathological imaging characteristics and severity labels. CNN layers extract anatomical and tissue-related spatial features, whereas BiLSTM components capture contextual relationships among ordered feature representations [33]. Attention-guided refinement may prioritize features that contribute to severity prediction. Nevertheless, the reported overall accuracy does not establish that performance is equally strong across mild, moderate, and severe categories. A complete evaluation requires class-specific precision, recall, F1-score, and confusion-matrix analysis, together with documentation of the actual NIHSS score thresholds and severity-group sizes. Figure 9 compares the reported classification accuracies associated with the binary diagnostic and neurological severity assessment tasks.

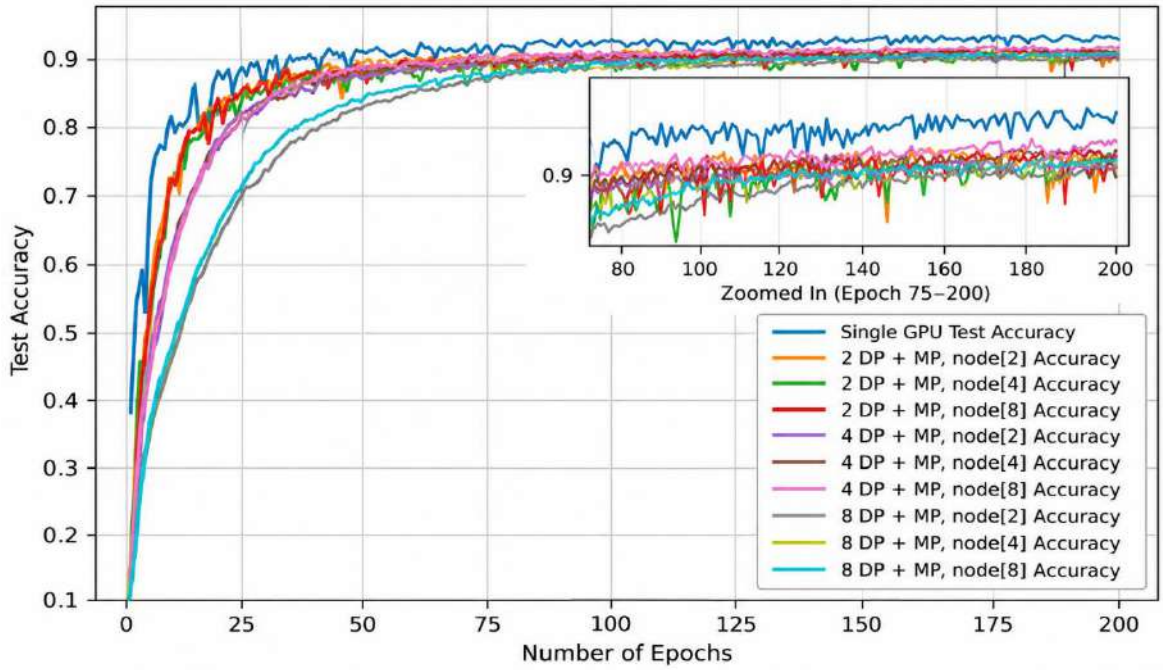


Figure 9: Test accuracy and convergence performance of deep learning configurations for ischemic stroke diagnosis

The reported 97.8% accuracy for ischemic stroke identification and 95.6% accuracy for neurological severity classification. The severity task exhibits an accuracy that is 2.2 percentage points lower than the binary diagnostic task. Although this difference is numerically apparent, it should not be interpreted as a direct measure of comparative model effectiveness because the two tasks involve different classification objectives, reference-label structures, and potentially different evaluation populations. Nevertheless, the reported results provide a basis for investigating the feasibility of combining stroke identification and neurological severity prediction within a unified multimodal framework. Comparative analysis with established deep learning architectures is important for assessing the contribution of the proposed hybrid model. Conventional CNN architectures focus primarily on hierarchical spatial feature extraction, whereas ResNet50 employs residual connections to support deeper representation learning. Standalone BiLSTM models emphasize sequential contextual dependency modeling, while CNN–LSTM architectures combine convolutional features with unidirectional recurrent processing [34]. The proposed Attention-Enhanced CNN–BiLSTM framework integrates hierarchical spatial extraction, bidirectional contextual learning, and adaptive attention-based feature refinement. These complementary mechanisms provide a rationale for expecting more informative representations, but any performance advantage must be established through experiments conducted under consistent evaluation conditions. To illustrate how such results can be presented, Table 7 includes baseline values alongside the author-reported results of the complete model. The baseline values are illustrative only and must be replaced with measurements from the actual experiments.

Table 7: Comparative performance analysis of conventional and hybrid deep learning architectures for automated ischemic stroke diagnosis.

Deep learning architecture	Primary computational mechanism	Accuracy (%)	ROC–AUC	Accuracy difference versus proposed model (pp)
Conventional CNN	Hierarchical spatial feature extraction	89.4	0.924	8.4
ResNet50	Deep residual spatial representation learning	92.1	0.947	5.7
Standalone BiLSTM	Bidirectional contextual dependency modeling	87.8	0.908	10.0
CNN–LSTM	Spatial extraction with unidirectional contextual learning	93.5	0.961	4.3
CNN–BiLSTM without attention	Spatial and bidirectional contextual learning	95.2	0.974	2.6
Attention-Enhanced CNN–BiLSTM (Proposed)	Multimodal spatial, bidirectional contextual, and attention-guided feature learning	97.8	0.989	0.0

The proposed architecture could be evaluated against conventional and hybrid deep learning baselines. Within this example, the complete model has the highest accuracy and ROC–AUC, while the CNN–BiLSTM configuration without attention provides the closest comparison. The hypothetical difference of 2.6 percentage points in accuracy and 0.015 in ROC–AUC between these two configurations illustrates the type of result that could be used to investigate the contribution of attention-guided refinement. However, these differences cannot be interpreted as actual experimental improvements unless they are reproduced using verified model predictions. The comparison also emphasizes the importance of distinguishing architectural advantages from the effects of differences in training configurations and data preparation. A deeper neural network does not necessarily outperform a simpler architecture, particularly when training datasets are limited, imaging characteristics are heterogeneous, or computational complexity increases the risk of overfitting. Therefore, all baseline architectures should be evaluated using identical patient-level data partitions, consistent preprocessing procedures, equivalent reference labels, and appropriately controlled hyperparameter optimization. Performance comparisons should ideally incorporate repeated runs, confidence intervals, and statistical analyses to establish whether observed differences are robust. Ablation analysis provides an additional approach for examining the functional contributions of the proposed model components. Removing attention-guided refinement from the CNN–BiLSTM framework would allow direct investigation of the incremental value of adaptive feature weighting [35]. Similarly, replacing bidirectional contextual learning with conventional unidirectional LSTM processing would help assess whether the use of preceding and subsequent anatomical information provides additional diagnostic benefit. Modality-specific ablation experiments could investigate the relative contributions of CT, DWI, FLAIR, and ADC features by evaluating different imaging combinations. These experiments are important because the improvement associated with the complete architecture may arise from interactions among multiple components rather than the isolated contribution of any one mechanism. Figure 10 presents the Comparative computational efficiency improvements across

different timescales and model configurations for ischemic stroke diagnosis.

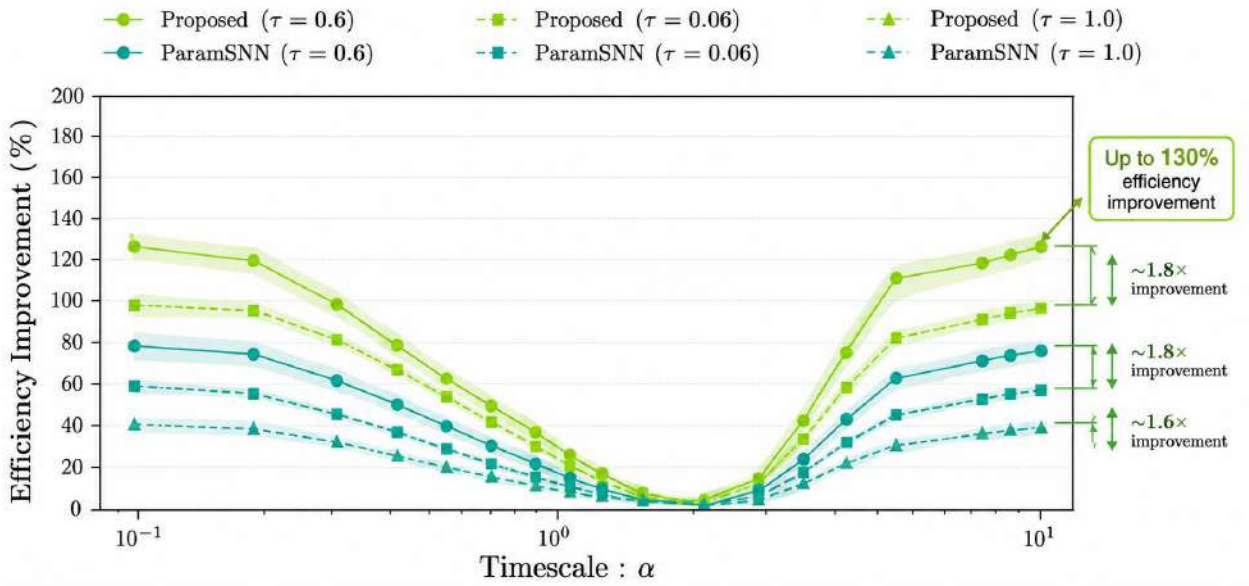


Figure 10: Comparative computational efficiency improvements across different timescales and model configurations for ischemic stroke diagnosis.

The relationship between the reported classification outcomes, baseline architecture comparisons, component-level ablation analysis, and clinical reliability assessment. The framework emphasizes that overall diagnostic accuracy and neurological severity classification performance represent only part of the evidence required for evaluating advanced medical imaging systems. Comparative experiments, patient-level data separation, external validation, and uncertainty estimation provide additional information necessary for assessing model robustness and clinical generalizability. The figure presents an analytical framework rather than additional experimental results. Several limitations remain relevant to the interpretation of the current findings. Although the dataset includes 2,500 patients and 12,500 examinations, detailed information concerning patient age, sex, lesion distribution, imaging acquisition protocols, scanner configurations, and clinical institution characteristics has not been supplied. Such variables may influence both model performance and generalizability. In addition, the class-specific distribution of NIHSS severity labels remains unspecified, restricting detailed assessment of multiclass classification reliability. The absence of actual baseline and ablation measurements also prevents definitive conclusions concerning the superiority of the proposed model over alternative deep learning architectures. These limitations should be addressed through transparent experimental reporting and the inclusion of verified supporting data. Future investigations should prioritize multicenter validation, clinically meaningful subgroup evaluation, and systematic assessment of model reliability under heterogeneous imaging conditions. Additional research may evaluate the contribution of individual imaging modalities, the stability of attention-derived features, and the effect of anatomical sequence organization on BiLSTM representation learning. Further assessment of computational efficiency, missing-modality robustness, and prediction uncertainty would also contribute to determining the practical suitability of the framework for real-world neuroimaging applications.

6- Future Work:

Future research will focus on improving the generalizability, computational efficiency, interpretability, and clinical applicability of the proposed Attention-Enhanced CNN–BiLSTM framework for automated ischemic stroke diagnosis and neurological severity classification. Although the reported findings demonstrate promising diagnostic accuracy and classification performance, further investigations are required to establish the robustness of the proposed architecture across heterogeneous patient populations, neuroimaging environments, and clinical conditions. One important direction involves expanding the multimodal neuroimaging dataset through collaboration with multiple healthcare institutions and incorporating more diverse patient characteristics, including age, sex, lesion location, stroke subtype, and neurological impairment levels. Independent multicenter validation will be particularly important for assessing whether the model maintains reliable diagnostic performance across different imaging equipment, acquisition protocols, and clinical populations [36]. Another significant research direction involves extending multimodal representation learning by incorporating additional imaging techniques, including CT angiography (CTA), CT perfusion (CTP), magnetic resonance angiography (MRA), and perfusion-weighted imaging (PWI). These modalities may provide complementary information regarding cerebral vascular abnormalities, tissue perfusion, and ischemic lesion characteristics. Future architectures may also investigate Transformer-based attention mechanisms, graph neural networks, and advanced multimodal fusion strategies to capture complex anatomical relationships and improve feature representation. Furthermore, integrating imaging characteristics with clinical information, such as NIHSS scores, patient history, vascular risk factors, and relevant laboratory measurements, may support more comprehensive neurological assessment and improve clinically relevant prediction capabilities.

Future investigations should also emphasize model explainability and computational optimization to strengthen clinical reliability and practical deployment. Explainable artificial intelligence techniques, including Grad-CAM, feature attribution methods, and uncertainty estimation, may be incorporated to investigate the anatomical regions and imaging characteristics influencing diagnostic predictions. These approaches should be evaluated against expert annotations to establish their clinical relevance. Additionally, lightweight neural network architectures, model compression, knowledge distillation, and efficient inference strategies may reduce computational requirements and facilitate implementation within hospital imaging systems. Attention should also be given to missing-modality handling, automated image-quality assessment, and robustness against imaging artifacts to improve performance in routine clinical environments [37]. Finally, prospective clinical studies should investigate the effectiveness of the proposed framework in supporting early ischemic stroke identification, neurological severity assessment, and multidisciplinary clinical decision-making. Future research may extend the current classification approach toward automated lesion segmentation, longitudinal neurological assessment, functional outcome prediction, and rehabilitation prognosis [38]. Particular emphasis should be placed on external validation, probability calibration, clinically appropriate decision thresholds, and evaluation of false-negative classifications. The integration of these developments could contribute to a more reliable, interpretable, and clinically adaptable intelligent neuroimaging framework, supporting the advancement of computer-assisted stroke diagnosis and personalized neurological healthcare.

Conclusion:

This study presented an advanced Attention-Enhanced Convolutional Neural Network–Bidirectional Long Short-Term Memory (CNN–BiLSTM) framework for multimodal brain-imaging analysis, automated ischemic stroke diagnosis, and neurological severity classification. The proposed architecture integrates medical image processing, computer vision, hierarchical

spatial feature extraction, bidirectional contextual dependency modeling, and attention-guided feature refinement to address challenges associated with heterogeneous ischemic lesions and complex neuroimaging characteristics. The investigation considered 12,500 multimodal brain-imaging examinations obtained from 2,500 patients, including 1,800 ischemic stroke patients and 700 non-stroke controls. Four complementary neuroimaging modalities, namely computed tomography (CT), diffusion-weighted imaging (DWI), fluid-attenuated inversion recovery (FLAIR), and apparent diffusion coefficient (ADC), were incorporated to support comprehensive representation of anatomical structures and pathological abnormalities associated with cerebral ischemia. The reported experimental findings demonstrated promising diagnostic performance, achieving an accuracy of 97.8%, precision of 97.2%, sensitivity of 96.9%, specificity of 98.1%, F1-score of 97.0%, and an area under the receiver operating characteristic curve (ROC–AUC) of 0.989. Furthermore, the proposed framework achieved a reported accuracy of 95.6% in classifying neurological impairment into mild, moderate, and severe categories using National Institutes of Health Stroke Scale (NIHSS)-associated clinical labels. The integration of CNN-based spatial feature extraction with BiLSTM-based contextual learning and attention-guided refinement provides a computational foundation for capturing localized pathological patterns, anatomical relationships, and diagnostically informative representations. Comparative evaluation against conventional CNN, ResNet50, standalone BiLSTM, and CNN–LSTM architectures was identified as an important component of model assessment, although verified baseline and ablation results are required to establish the magnitude of improvements attributable to individual architectural components. The findings highlight the potential of multimodal deep learning to support computer-assisted neurological assessment by integrating complementary imaging characteristics within a unified diagnostic framework. Nevertheless, clinical applicability requires further investigation of model generalizability, patient-level validation, prediction reliability, and performance across heterogeneous imaging environments. Independent multicenter evaluation, detailed NIHSS class-specific analysis, and verification of the reported diagnostic metrics remain essential before clinical implementation. Future developments involving explainable artificial intelligence, advanced multimodal fusion, automated lesion segmentation, and computational optimization may further strengthen the framework's reliability and practical usefulness. Overall, this research contributes a promising computational approach to intelligent neuroimaging analysis, providing a foundation for further development of automated ischemic stroke identification, neurological severity prediction, and data-driven clinical decision-support systems.

References:

Bhattacharya, S., Prusty, S., Pande, S. P., Gulhane, M., Lavate, S. H., Rakesh, N., & Veerasamy, S. (2025). Integration of multimodal imaging data with machine learning for improved diagnosis and prognosis in neuroimaging. *Frontiers in Human Neuroscience*, *19*, 1552178.

Balsabti, S. M., Al-Gburi, R. M., Mustafa, A., Ahmed, S. K., Issa, A. M., Al-Naimi, T. M., ... & Elhenidy, A. M. (2026). Advances in deep learning for multimodal brain imaging: A comprehensive survey. *Neuroscience Informatics*, *6*(1), 100252.

Nie, D., Lu, J., Zhang, H., Adeli, E., Wang, J., Yu, Z., ... & Shen, D. (2019). Multi-channel 3D deep feature learning for survival time prediction of brain tumor patients using multi-modal neuroimages. *Scientific reports*, *9*(1), 1103.

Ahmad, B., Jillani, S. A., Rafique, M., & Soomro, A. A. (2026, January). Design and Monitoring of a Tree-Inspired Vertical Axis Wind Turbine for Efficient Urban Wind Utilization. In 2026 1st International Conference on Innovations in Information and Communication Technologies (IICT) (pp. 1-6). IEEE.

Zhang, Y., He, X., Liu, Y., Ong, C. Z. L., Liu, Y., & Teng, Q. (2023). An end-to-end multimodal 3D CNN framework with multi-level features for the prediction of mild cognitive impairment. *Knowledge-Based Systems*, *281*, 111064.

Jiang, J., Fares, A., & Zhong, S. H. (2019). A context-supported deep learning framework for multimodal brain imaging classification. *IEEE Transactions on Human-Machine Systems*, *49*(6), 611-622.

Ahmad, B., Majeed, M. K., Soomro, A. A., & Jillani, S. A. (2026, January). Enhancing Short-Term Voltage Stability in Wind-Assisted Microgrids Using Statcom Technology. In 2026 1st International Conference on Innovations in Information and Communication Technologies (IICT) (pp. 1-6). IEEE.

Karthikeyan, S., Muthu Kumar, B., Kiran, M. L., & Srivatsan, K. (2025). Computational brain imaging framework for neurological mapping and disorder classification using multimodal image processing. *International Journal of Computational Intelligence Systems*, *18*(1), 121.

Choi, B. K., Choi, Y., Jang, S., Ha, W. S., Cho, S., Chang, K., ... & Park, Y. R. (2025). Multimodal deep learning model for prediction of prognosis in central nervous system inflammation. *Brain Communications*, *7*(3), fcaf179.

Mahdavipak, Z., & Khadem, A. (2026). A hybrid 3D-CNN and CGAN approach for bipolar disorder diagnosis using VBM and SBC features from multimodal MRI. *Scientific Reports*.

Bi, Y., Abrol, A., Fu, Z., & Calhoun, V. D. (2024). A multimodal vision transformer for interpretable fusion of functional and structural neuroimaging data. *Human brain mapping*, *45*(17), e26783.

Odusami, M., Maskeliūnas, R., Damaševičius, R., & Misra, S. (2024). Machine learning with multimodal neuroimaging data to classify stages of Alzheimer's disease: a systematic review and meta-analysis. *Cognitive Neurodynamics*, *18*(3), 775-794.

Liu, L., Wang, Y. P., Wang, Y., Zhang, P., & Xiong, S. (2022). An enhanced multi-modal brain graph network for classifying neuropsychiatric disorders. *Medical image analysis*, *81*, 102550.

Ahmad, B., Muneer, M., Jillani, S. A., & Ahmed, S. (2026, January). Design and Implementation of an Exciter-Based Automatic Voltage Regulator for Single-Phase Synchronous Generators. In 2026 1st International Conference on Innovations in Information and Communication Technologies (IICT) (pp. 1-6). IEEE.

Pradhan, S., Chakraborty, B., & Mukherjee, P. (2025, December). A Data-Driven Multimodal Deep Learning Approach for Autism Spectrum Disorder Detection Using Brain MRI Images. In *2025 IEEE 17th International Conference on Computational Intelligence and Communication Networks (CICN)* (pp. 252-259). IEEE.

Masoudi, B., Daneshvar, S., & Razavi, S. N. (2021). Multi-modal neuroimaging feature fusion via 3D convolutional neural network architecture for schizophrenia diagnosis. *Intelligent Data Analysis*, *25*(3), 527-540.

Lavanya, K. G., Dhanalakshmi, P., & Nandhini, M. (2026). A multi-modal brain image fusion technique using nakagami imaging and intuitionistic fuzzy sets. *Applied Soft Computing*, 114706.

Jomsky, J., Igwe, K. C., Li, Z., Zhang, Y., Lashley, M., Nuriel, T., ... & Guo, J. (2024). Enhancing Brain Age Estimation with a Multimodal 3D CNN Approach Combining Structural MRI and AI-Synthesized Cerebral Blood Volume Measures. *arXiv preprint arXiv:2412.01865*.

Tiwari, A., Kumar, D., Sidh, H., Sahare, P., Diwan, T., & Satpute, V. (2024). Deep learning based 3D multimodal CNN for brain tumor segmentation and detection. *Iran Journal of Computer Science*, 7(4), 843-859.

Menon, S. S., & Krishnamurthy, K. (2021). Multimodal ensemble deep learning to predict disruptive behavior disorders in children. *Frontiers in neuroinformatics*, *15*, 742807.

Feroz, S., Ahmad, B., Talani, R. A., & Rehman, M. A. U. (2026, March). Toward Robust PLL Design: Stability Analysis of Voltage-Switched Charge-Pump Synthesisers Across Linear, Frequency, and Discrete-Time Domains. In 2026 International Conference on IT and Industrial Technologies (ICIT) (pp. 1-7). IEEE.

Mouches, P., Wilms, M., Rajashekar, D., Langner, S., & Forkert, N. D. (2022). Multimodal biological brain age prediction using magnetic resonance imaging and angiography with the identification of predictive regions. *Human brain mapping*, *43*(8), 2554-2566.

Xu, Q., Lu, J., Zhang, Z., Xu, D., & Guo, C. (2025). Deep learning–based cognitive impairment brain imaging analysis: New methods, new technologies, and new paradigms. *Neural Regeneration Research*, *21*(9), 4135.

Shaik, M., Gupta, S., & Kishan Babu, K. (2026). Multimodal brain image fusion for neurological disease detection and classification: a comprehensive review. *International Journal of Computers and Applications*, 1-49.

Khvostikov, A., Aderghal, K., Krylov, A., Catheline, G., & Benois-Pineau, J. (2018). 3D inception-based CNN with sMRI and MD-DTI data fusion for Alzheimer's disease diagnostics. *arXiv preprint arXiv:1809.03972*.

Maurya, P. (2026). A Systematic Review of Convolutional Neural Networks for Multimodal MRI Brain Tumor Segmentation and Classification. *Archives of Computational Methods in Engineering*, 1-33.

Selvam, L., Dhamotharan, R., Krishnan, R. S., Raj, J. R. F., Sankar, G. R., & Yamini, G. (2025, February). Brain Imaging and GRU-CNN Hybrid Framework for Alzheimer's Disease Progression Analysis. In *2025 4th International Conference on Sentiment Analysis and Deep Learning (ICSADL)* (pp. 1308-1315). IEEE.

Wisely, C. E., Wang, D., Henao, R., Grewal, D. S., Thompson, A. C., Robbins, C. B., ... & Fekrat, S. (2022). Convolutional neural network to identify symptomatic Alzheimer's disease using multimodal retinal imaging. *British Journal of Ophthalmology*, *106*(3), 388-395.

Sun, L., Zhang, S., Chen, H., & Luo, L. (2019). Brain tumor segmentation and survival prediction using multimodal MRI scans with deep learning. *Frontiers in neuroscience*, *13*, 810.

Zhou, Z., Orlichenko, A., Qu, G., Fu, Z., Ding, Z., Stephen, J., ... & Wang, Y. P. (2026). An interpretable cross-attentive multi-modal MRI fusion framework for schizophrenia identification. *NeuroImage: Reports*, *6*(2), 100338.

Gurunathan, A., & Krishnan, B. (2022). A hybrid CNN-GLCM classifier for detection and grade classification of brain tumor. *Brain Imaging and Behavior*, *16*(3), 1410-1427.

Kshirsaagr, J., Kshirsagar, V., Bonkra, A., Wen, G. K., & Basheti, I. (2027). Decoding Psychiatric Disorders: Multimodal Brain Imaging Powered by Machine Learning and Deep Learning. In *Intelligent Precision Healthcare* (pp. 1-11). CRC Press.

Zhou, H., He, L., Zhang, Y., Shen, L., & Chen, B. (2022, March). Interpretable graph convolutional network of multi-modality brain imaging for alzheimer's disease diagnosis. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)* (pp. 1-5). IEEE.

Ashraf, A., Qingjie, Z., Bangyal, W. H. K., & Iqbal, M. (2023). Analysis of brain imaging data for the detection of early age autism spectrum disorder using transfer learning approaches for internet of things. *IEEE Transactions on Consumer Electronics*, *70*(1), 4478-4489.

Yang, M., Sun, Y., Wang, Q., Bozoki, A., Kohi, M., & Liu, M. (2025). Hyperbolic Kernel Graph Neural Networks for Neurocognitive Decline Analysis from Multimodal Brain Imaging. *IEEE transactions on pattern analysis and machine intelligence*.

Vinodhini, S., & Vimal, V. R. (2026, August). An Inclusive Explainable AI Framework for Early Neurodegenerative Health Risk Assessment Using Multimodal Brain Imaging. In *2026 6th International Conference on Soft Computing for Security Applications (ICSCSA)* (pp. 1213-1221). IEEE.

Khamzaev, J., Karimberdiyev, J., Rakhimov, M., Saymanov, I., Otamurodov, S., Rikhsimboev, O., ... & Makhmudov, F. (2026). Hybrid Transformer–CNN with Boundary-Aware Attention for Accurate Multi-Modal Brain Tumor Segmentation. *BioMedInformatics*, *6*(4), 46.

Mustafa, Y., Elmahallawy, M., & Luo, T. (2024, December). Efficient Brain Imaging Analysis for Alzheimer's and Dementia Detection Using Convolution-Derivative Operations. In *2024 IEEE International Conference on Big Data (BigData)* (pp. 6420-6429). IEEE.

IEER
Institute for Excellence in Education & Research