

Multimodal Deep Learning and Transcriptomic Synthesis for Precision Diagnostics in Acute Ischemic Stroke

***J. Jackulin Reeja, V. S. Harilakshmi, F Fanax Femy, J. Anto Hepzie Bai,
B. S. Saravanabala and Nitha Justin**

Department of Computer Science, Holy Cross College (Autonomous), Nagercoil - 629004

Affiliated to Manonmaniam Sundaranar University, Tirunelveli - 627012

*Corresponding Author - Email: jackulinreeja@holycrossnagl.edu.in

ABSTRACT

Acute ischemic stroke (AIS) is a major cause of mortality and long-term neurological disability, requiring rapid and accurate diagnosis for effective treatment. This study proposes a Hybrid Attention Fusion Multimodal Deep Learning (HAF-MDL) framework that integrates brain magnetic resonance imaging (MRI) with whole-blood transcriptomic biomarkers for precision diagnosis of AIS. DWI, ADC, and FLAIR MRI sequences are analyzed using a ResNet50-UNet with a Convolutional Block Attention Module (CBAM) for automated lesion segmentation. Transcriptomic data are processed through differential gene expression analysis, WGCNA, PPI network analysis, Boruta, LASSO, and SVM-RFE, identifying six key biomarkers: ATG7, KAT2A, RNF20, UBA1, UBE2I, and USP15. An Artificial Neural Network (ANN) classifier is developed using these biomarkers, and intermediate feature fusion combines imaging and molecular representations for unified diagnosis and prognosis prediction. Experimental results demonstrate that the proposed HAF-MDL framework outperforms unimodal approaches, achieving higher diagnostic sensitivity, specificity, and AUC. The proposed multimodal framework provides an interpretable and clinically applicable approach for accurate AIS diagnosis, prognostic assessment, and personalized treatment planning, supporting precision medicine in stroke management.

Keywords: *Acute ischemic stroke; Deep learning; Magnetic resonance imaging (MRI); Transcriptomics; Multimodal learning; Biomarkers.*

1. Introduction

Historically, acute ischemic stroke diagnosis has relied on rapid clinical evaluations and neuroimaging modalities. Non-contrast computed tomography is widely employed due to its accessibility and speed, yet it exhibits limited sensitivity in identifying early hyper-acute infarctions. Contrastingly, magnetic resonance imaging provides unparalleled soft-tissue contrast, high spatial resolution, and exquisite sensitivity to early microstructural changes [1]. However, manual evaluation of multi-slice magnetic resonance imaging volumes is time-consuming, subjective, and highly dependent on operator experience, which can introduce cognitive delays and inter-rater inconsistencies in high-stakes emergency settings.

To overcome these barriers, a paradigm shift is occurring toward automated computer-aided diagnosis systems leveraging deep learning [2]. By training deep convolutional neural networks on multi-parametric magnetic resonance imaging sequences, clinical systems can rapidly segment ischemic lesions and classify stroke territories. Yet, macro-scale imaging alone cannot fully capture the dynamic cellular and molecular microenvironment driving tissue fate. The next frontier in precision stroke medicine relies on the fusion of macro-scale neuroimaging with micro-scale molecular data, specifically whole-blood transcriptomics reflecting systemic post-translational modifications. This multimodal synthesis allows for a holistic assessment of both anatomical injury and physiological viability, providing a robust pathway for early risk stratification, prediction of long-term functional recovery, and individualized therapeutic selection.

1.1 High-Resolution Neuroimaging Modalities and Ischemic Markers

Magnetic resonance imaging acts as the clinical gold standard for characterizing brain ischemia, utilizing several core sequences that capture distinct aspects of tissue pathophysiology. Diffusion-Weighted Imaging (DWI) is uniquely sensitive to the random microscopic motion of water molecules. Within minutes of arterial occlusion, the failure of energy-dependent sodium-potassium pumps leads to an influx of extracellular water into the intracellular compartment, causing cytotoxic edema. This cellular swelling restricts water diffusion, rendering the acute infarct core highly hyperintense on diffusion-weighted images [3,4].

To verify true diffusion restriction and eliminate confounding "T2 shine-through" effects, diffusion-weighted imaging is always paired with Apparent Diffusion Coefficient (ADC) maps. True acute ischemic lesions present as distinct hypointensities on apparent diffusion coefficient maps, reflecting a drop in the quantitative diffusion coefficient. Fluid-Attenuated Inversion Recovery (FLAIR) suppresses free fluid (cerebrospinal fluid) signals, making it ideal for identifying vasogenic edema, subacute tissue damage, and pre-existing cerebral small vessel disease [5].

Integrating these sequences provides critical biological markers that guide therapeutic decisions. The DWI-FLAIR mismatch represents a crucial clinical surrogate for stroke onset timing. Because hyper-acute ischemic lesions (less than 4.5 hours from onset) appear bright on diffusion-weighted imaging but are not yet visible on fluid-attenuated inversion recovery sequences, the presence of a mismatch indicates that the patient is within the safe window for intravenous thrombolysis. Once the lesion becomes hyperintense on fluid-attenuated inversion recovery, it denotes substantial vasogenic edema, indicating a subacute phase and

an elevated risk of reperfusion-induced hemorrhage. Additionally, structural markers such as total lesion volume, spatial location, and overlap with critical eloquent motor pathways—most notably the corticospinal tract—serve as strong, interpretable predictors of long-term motor recovery and functional disability, typically quantified via the modified Rankin Scale [6,7,8].

1.2 Systemic Transcriptomic Biomarker Extraction

Computational analysis of peripheral whole-blood transcriptomic datasets has emerged as a non-invasive, translationally powerful strategy to uncover molecular biomarkers of acute ischemia [9]. Whole-blood transcriptomic datasets, specifically GSE16561 (comprising 39 ischemic stroke patients and 24 healthy controls, total n=63) and GSE58294 (comprising 69 cardiogenic embolic stroke patients and 23 controls, total n=92), have been analyzed using robust bioinformatics workflows. Low-expression filtering (levels ≤ 1 in >50% of samples), K-nearest neighbors (KNN) imputation, quantile normalization, and $\log_2(\text{expression} + 1)$ transformation are performed on the microarrays to establish consistent baseline distributions.

Weighted Gene Co-expression Network Analysis (WGCNA) using a soft power threshold of 7 identifies functional gene modules, revealing a significant positive correlation between the brown module and ischemic stroke, and negative correlations for the black, magenta, and tan modules. Differentially expressed genes (DEGs) are identified using empirical Bayes linear modeling with criteria of $|\log_2\text{Fold Change}| \geq 0.2$ and adjusted $P < 0.05$, yielding 1,465 upregulated and 1,782 downregulated genes. The intersection of significant WGCNA module genes, DEGs, and known post-translational modification-related genes (PTMRGs) identifies 75 candidate biomarkers.

A Protein-Protein Interaction (PPI) network topology (using the STRING database with a confidence threshold >0.4 and Cytoscape's cy to Hubba plugin) prioritizes the top 20 connected nodes. Three machine learning feature selection methods are then applied to these nodes: Boruta (confirming 13 features via Bonferroni correction at $P < 0.01$), LASSO regression (yielding 10 core features through L1 regularization), and SVM-RFE (identifying 10 features via 10-fold cross-validation), intersecting to yield 7 robust biomarkers. Removing the gene NEDD8 due to inconsistent external expression trends yields six highly prioritized post-translational modification-related diagnostic biomarkers. Following Table 1 consists of gene expression changes and pathophysiological impact in ischemia.

Table 1 Gene Expression Changes and Pathophysiological Impact in Ischemia

Gene Symbol	Full Biological Name	Direction of Change	Primary Post-Translational Modification Role	Downstream Pathophysiological Contribution in Ischemia
ATG7	Autophagy Related 7	Upregulated	Autophagic conjugation enzyme	Initiates autophagosome membrane formation; over-activation triggers cell death.
KAT2A	Lysine Acetyltransferase 2A	Downregulated	Histone/non-histone acetylation	Controls chromatin remodeling; depletion induces anti-inflammatory microglial shift.
RNF20	Ring Finger Protein 20	Upregulated	E3 ubiquitin-protein ligase	Monoubiquitinates Histone H2B; cooperatively drives STAT3 astrocytic differentiation.
UBA1	Ubiquitin Activating Enzyme 1	Upregulated	Primary ubiquitin activating enzyme	Initiates cellular ubiquitination cascade; crucial for proteotoxic clearance.
UBE2I	Ubiquitin Conjugating Enzyme E2 I	Downregulated	SUMOylation conjugating enzyme	Serves as the sole conjugator for SUMOylation-induced cell survival mechanisms.
USP15	Ubiquitin Specific Peptidase 15	Upregulated	Deubiquitinating cysteine protease	Removes ubiquitin chains; regulates type I interferon pathway and neuroinflammation.

2. Literature Review

Magnetic Resonance Imaging (MRI) remains the preferred imaging modality for evaluating acute ischemic stroke because of its superior soft-tissue contrast and high sensitivity in detecting early ischemic lesions. Diffusion-Weighted Imaging (DWI) identifies cytotoxic edema within minutes after arterial occlusion, whereas Apparent Diffusion Coefficient (ADC) maps differentiate true diffusion restriction from T2 shine-through artifacts. Fluid-Attenuated Inversion Recovery (FLAIR) imaging provides complementary information regarding lesion age and tissue viability through the clinically important DWI–FLAIR mismatch [10,11].

Traditional computer-aided diagnosis systems relied on handcrafted image features combined with classical machine learning classifiers such as Support Vector Machines (SVM), Random Forests, and K-Nearest Neighbors (KNN). Although these methods achieved moderate diagnostic performance, they required manual feature engineering and often failed to capture complex lesion morphology, resulting in poor generalization across different patient populations and imaging protocols [12].

The rapid advancement of deep learning has significantly improved automated stroke diagnosis. Fully Convolutional Networks (FCNs), U-Net, ResNet, DenseNet, EfficientNet, and Vision Transformer (ViT)-based architectures have demonstrated superior performance in brain lesion segmentation by learning hierarchical image representations directly from MRI data [13]. Among these, U-Net has become one of the most widely adopted architectures for biomedical image segmentation due to its encoder–decoder structure and skip connections that preserve spatial information during feature reconstruction [14].

3. Methodology

3.1 Overview of the Proposed Framework

The proposed framework integrates complementary structural information extracted from multi-parametric MRI, with molecular information obtained from whole-blood transcriptomic biomarkers. Unlike conventional unimodal diagnostic systems that rely solely on imaging or genomic information, the proposed multimodal architecture simultaneously captures anatomical abnormalities and systemic molecular responses associated with cerebral ischemia. Brain MRI data are obtained from publicly available ischemic stroke datasets, including the ATLAS and ISLES repositories, which contain expert-annotated ischemic lesion masks for supervised deep learning. The complete workflow of the proposed framework is illustrated in Figure 1.

3.2 Artificial Neural Network Construction for Transcriptomic Classification

An Artificial Neural Network (ANN) is constructed to map the expression values of the six prioritized diagnostic biomarkers into a binary classification indicating ischemic stroke versus healthy control [15].




Fig. 1 Multi-model neural network architecture

The network is developed using the R package neuralnet. Prior to training, batch correction and Z-score normalization are applied across the datasets to align their feature distributions.

$$\text{Z-score} = \frac{x - \mu}{\sigma} \quad (1)$$

The network topology consists of an input layer accepting the six diagnostic gene expression features, a single hidden layer composed of 5 hidden neurons, and a single output neuron utilizing a logistic sigmoid activation function. The model is trained with a propagation error threshold set to 0.05 and a maximum constraint of 1×10^6 training steps to prevent infinite loop execution during gradient optimization. Rigorous 10-fold cross-validation is implemented during model development to minimize the risk of overfitting on the primary dataset.

$$\text{Sigmoid}(z) = \frac{1}{1 + e^{-z}} \quad (2)$$

To validate the clinical applicability of these six diagnostic markers, reverse transcription quantitative PCR (RT-qPCR) is performed on whole-blood samples from an independent validation cohort of 10 ischemic stroke patients and 10 healthy controls. Total RNA is isolated using TRIzol reagent, reverse-transcribed into cDNA, and amplified through 40 cycles. The relative expression levels are calculated using the comparative $2^{-\Delta\Delta C_t}$ method, with statistical significance determined via two-tailed student's t-tests ($P < 0.05$). This clinical validation confirms that ATG7, RNF20, UBA1, and USP15 are significantly upregulated in stroke patients, while KAT2A and UBE2I are significantly downregulated, matching the bioinformatics predictions. The comparative evaluation of diagnostic models listed in Table 2.

Table 2 Comparative Evaluation of Diagnostic Models

Diagnostic Model Configuration	Training Dataset AUC (GSE16561)	Testing Dataset AUC (GSE58294)	Clinical Sensitivity / Specificity Implications
ATG7 Standalone	0.750	0.934	Strong performance in subacute settings; reflects active autophagic stress.
KAT2A Standalone	0.853	0.723	Moderately sensitive; reflects early epigenetic transcription alterations.
RNF20 Standalone	0.864	0.845	Consistent diagnostic capacity; tracks chromatin changes and astrocyte activation.
UBA1 Standalone	0.702	0.749	Acceptable performance; reflects general proteasomal stress response.

UBE2I Standalone	0.848	0.736	Strong training AUC; marks the extent of microvascular SUMOylation.
USP15 Standalone	0.871	0.910	High performance; correlates with systemic immune response and neuroinflammation.
Six-Gene Integrated ANN	0.983	0.950	Superior performance; leverages non-linear correlations across pathways.

3.3 Computerized Architectures for Brain MRI Segmentation

Hybrid ResNet50-UNet Model (ResUNet50)

The automated segmentation of ischemic lesions in brain magnetic resonance imaging is formulated as a semantic pixel-wise classification task. While standard U-Net models perform reasonably well, they often struggle with boundary delineation of irregular, scattered, or microvascular infarctions. To address this, a hybrid ResNet50-UNet architecture replaces the standard contracting path with a pre-trained ResNet50 encoder. The ResNet50 backbone (pre-trained on ImageNet, with classification layers excluded using `include_top=False`) extracts high-level semantic features while bypassing vanishing gradient limitations through internal skip connections.

In a Keras and TensorFlow implementation, intermediate activations are captured from the encoder to serve as skip connections to the decoder. The exact Keras layer query names used to extract these skip feature maps are `input_1` (for raw resolution features), `conv1_relu`, `conv2_block3_out`, and `conv3_block4_out`.

The deepest layer of the ResNet50 encoder, `conv4_block6_out`, serves as the semantically rich, low-resolution bottleneck. Each decoder block upsamples the feature maps using transposed 2D convolutions, concatenates them with the corresponding encoder skip activations using Keras Concatenate layers, and refines the concatenated maps through consecutive 3×3 convolutional layers, batch normalization, and rectified linear unit (ReLU) activations.

Spatial and Channel Attention Integration

To enhance spatial awareness and minimize false-positive segmentations, a Convolutional Block Attention Module (CBAM) is integrated into the ResNet50-UNet framework. The CBAM block dynamically refines intermediate features by sequentially computing channel and spatial attention maps [11,12]. The channel attention submodule aggregates spatial details of a feature map using both Global Average Pooling (GAP) and Global Max Pooling (GMP) [13, 14]. These pooled features are processed through a shared

multi-layer perceptron to generate a channel-wise attention vector, highlighting "what" semantic features are most important.

The spatial attention submodule takes the channel-refined feature map, applies average and max-pooling operations along the channel axis, and convolving them with a 7×7 kernel to generate a spatial-wise attention map, highlighting "where" the ischemic regions are located. This joint attention gating mechanism suppresses noise from surrounding healthy tissues and artifacts, improving the model's focus on acute ischemic stroke lesions.

4. Experimental Setup

4.1 Model Compilation, Optimization, and Training Dynamics

To train deep learning models on brain magnetic resonance imaging datasets (e.g., ISLES or ATLAS challenges), specific training pipelines must be configured in Keras and TensorFlow to handle class imbalances, as stroke lesions typically occupy less than 2% of total brain voxels.

$$\text{Class Ratio} = \frac{\text{Lesion Voxels}}{\text{Total Brain Voxels}} \leq 0.02 \quad (3)$$

Advanced Combined Loss Formulations

To mitigate the impact of class imbalance, standard binary cross-entropy is combined with the Dice Similarity Coefficient (DSC) loss to form a hybrid loss function:

$$\mathcal{L}_{\text{Combined}} = \alpha \mathcal{L}_{\text{BCE}} + \beta \mathcal{L}_{\text{Dice}} \quad (4)$$

Binary cross-entropy computes the average voxel-wise classification error:

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (5)$$

The Dice loss is derived from the overlap between the predicted mask and the ground-truth label:

$$\mathcal{L}_{\text{Dice}} = 1 - \frac{2 \sum_{i=1}^N y_i \hat{y}_i + \epsilon}{\sum_{i=1}^N y_i + \sum_{i=1}^N \hat{y}_i + \epsilon} \quad (6)$$

where $\epsilon = 1 \times 10^{-7}$ is a smoothing factor to prevent division-by-zero errors. To further minimize false-negative predictions (which can lead to untreated penumbral tissue), a Tversky loss is used to adjust the penalties for false positives and false negatives:

$$\mathcal{L}_{\text{Tversky}} = 1 - \frac{\sum y_i \hat{y}_i + \epsilon}{\sum y_i \hat{y}_i + \gamma \sum (1 - y_i) \hat{y}_i + \delta \sum y_i (1 - \hat{y}_i) + \epsilon} \quad (7)$$

Setting $\delta = 0.7$ and $\gamma = 0.3$ penalizes false negatives more heavily, encouraging the model to maintain sensitivity to subtle lesion boundaries.

Optimization Strategy and Learning Rate Schedulers

The model is compiled using the Adam optimizer with an initial learning rate of 1×10^{-4} . A cosine decay learning rate scheduler is implemented to gradually reduce the learning rate to 1×10^{-6} over 150 epochs, facilitating stable convergence as training progresses:

$$\eta_t = \eta_{\min} + \frac{1}{2}(\eta_{\max} - \eta_{\min}) \left(1 + \cos \left(\frac{T_{\text{cur}}}{T_{\text{max}}} \pi \right) \right) \quad (8)$$

Model checkpointing preserves the weights that minimize validation loss, and early stopping with a patience of 15 epochs prevents overfitting on small medical datasets.

Data Augmentation Operations in Keras

To enhance model generalizability and prevent overfitting, data augmentation is performed on-the-fly during training using Keras pre-processing layers or the tf.data pipeline. The spatial transformations applied to both the input images and their corresponding segmentation masks include:

- Random rotation of the image's X-axis up to 15 degrees.
- Horizontal and vertical random translations up to 20% of image dimensions.
- Random image shearing within 10 degrees.
- Zoom factor variations up to 20%.
- Random brightness adjustments between 0.8 and 1.2.

$$\text{Augmentation Pipeline} = \{\text{Rotation}_{15^\circ}, \text{Translation}_{20\%}, \text{Shear}_{10^\circ}, \text{Zoom}_{20\%}\} \quad (9)$$

To maximize hardware utilization, all mathematical operations are wrapped in compiled graph functions using the @tf.function decorator, and the model leverage mixed precision training (float16 weights with a float32 loss accumulator) to accelerate training on compatible GPUs.

5. Results and Discussion

5.1 Multimodal Data Fusion and Systemic Prognostic Modeling

To overcome the diagnostic limitations of single-modality systems, intermediate (joint) feature fusion is used to integrate macroscopic neuroimaging with systemic genomic profiling. The conceptual formulation of intermediate fusion involves extracting latent features from each modality through independent, specialized network branches, and

combining them at a deeper hidden layer. This enables the model to preserve modality-specific details while learning non-linear cross-modal correlations during joint optimization.

$$\text{Intermediate Joint Feature Space: } M_{\text{Fused}} = M_{\text{MRI}} \oplus M_{\text{Genomic}} \quad (10)$$

The clinical outcome model is trained to predict a binary prognosis at discharge (modified Rankin Scale score ≥ 3 , indicating moderate-to-severe disability or death). The imaging branch processes multi-parametric MRI sequences (DWI, ADC, and FLAIR) to extract the total lesion volume, spatial coordinates, and the degree of corticospinal tract (CST) overlap. The genomic branch processes the six-dimensional transcriptomic expression vector (ATG7, KAT2A, RNF20, UBA1, UBE2I, USP15).

The fused representation is mapped through joint dense layers with dropout regularization (0.3) and a final sigmoid output node. This multimodal fusion approach consistently outperforms individual unimodal baselines, showing a substantial increase in diagnostic accuracy, prognostic sensitivity, and calibration stability. The Table 3 shows the diagnostic performance of multimodal deep learning architectures.

Table 3 Diagnostic Performance of Multimodal Deep Learning Architectures

Model Architecture	Input Data Modalities	Target Clinical Endpoint	Diagnostic Sensitivity	Diagnostic Specificity	Area Under the Curve (AUC)
Unimodal Imaging Model	Brain MRI (DWI/ADC)	3-Month m RS Prognosis	0.695	0.720	0.737
Unimodal Clinical Model	Patient Tabular Metadata	3-Month m RS Prognosis	0.708	0.878	0.808
Unimodal Genomic Model	Whole-Blood 6-Gene Panel	Acute Stroke Triage	0.739	0.902	0.950
Early Fusion Model	MRI + Genomic Vectors	3-Month m RS Prognosis	0.723	0.891	0.825
Late Fusion Model	Model Probability Averaging	3-Month m RS Prognosis	0.731	0.915	0.850
Intermediate Fusion Model	Joint Keras Functional API	3-Month m RS Prognosis	0.875	0.939	0.975
HAF-MDL Model	Image + Genomics Cross-Attention	Brain Pathology Triage	0.912	0.960	0.964

5.2 Biological Mechanisms and Computational Docking Validation

To validate the clinical relevance of the six post-translational modification-related diagnostic genes, Gene Set Enrichment Analysis (GSEA) was performed on the GSE16561

dataset. The analysis linked these biomarkers to critical biological pathways activated during cerebral ischemia:

- ATG7 is significantly enriched in 27 signaling pathways, including DNA repair, E2F target genes, and MYC-mediated transcription, which govern downstream autophagy and neuronal survival.
- KAT2A is enriched in 19 pathways, including the coagulation cascade, complement pathways, and heme metabolism, which influence early microvascular thrombosis.
- RNF20 relates to 25 pathways, including DNA damage repair, oxidative phosphorylation, and MYC signaling, promoting astrocyte differentiation via STAT3 epigenetic activation.
- UBA1 is enriched in 20 pathways, including epithelial-mesenchymal transition and E2F target genes, reflecting active proteological stress.

UBE2I and USP15 are enriched in pathways governing glycolysis, DNA repair, and MYC signaling, which modulate the cellular response to hypoxia. The Figure 2 illustrates the biological pathway and docking validation.




Fig. 2 Biological pathway and docking validation

Evaluating the immune cell microenvironment using the CIBERSORT algorithm reveals distinct immune cell infiltration patterns between ischemic stroke patients and controls. Significant differences are observed in peripheral CD8⁺ cytotoxic T cells ($P < 0.001$), activated CD4⁺ helper T cells ($P < 0.05$), M0 macrophages ($P < 0.01$), and neutrophils ($P < 0.0001$). Spearman correlation analysis indicates that:

- ATG7, RNF20, and UBA1 are significantly associated with M0 macrophage infiltration, reflecting systemic activation of phagocytic cells.
- KAT2A and UBE2I correlate with activated CD4⁺ and CD8⁺T cells, which can cross the blood-brain barrier and exacerbate tissue damage.
- USP15 expression correlates with peripheral neutrophil levels, which can release reactive oxygen species and contribute to blood-brain barrier breakdown.

To identify potential small-molecule therapies, the Enrichr and Drug Gene Interaction Database (DGIdb) were queried using the six diagnostic genes. Eicosatetraynoic acid (ETYA), a synthetic triple-bonded arachidonic acid analogue, was predicted as a promising therapeutic candidate. Molecular docking performed using CB-DOCK2 shows high binding affinity between ETYA and KAT2A (Vina score of -7.0 kcal/mol) and UBE2I (Vina score of -5.6 kcal/mol).

ETYA binds within the active pocket of KAT2A by forming two hydrogen bonds with the arginine residue (ARG) at position 45. It binds to UBE2I by forming a single hydrogen bond with the arginine residue at position 62. These findings suggest that ETYA may limit ischemic injury by modulating SUMOylation and histone acetylation pathways. The following Table 4 computational validation of protein–compound interactions through molecular docking studies.

Table 4 Computational Validation of Protein–Compound Interactions Through Molecular Docking Studies

Target Protein	UniProt ID	Interacting Compound	PubChem CID	Vina Score (kcal/mol)	Active Docking Pocket Binding Site residues
KAT2A	Q92830	Eicosatetraynoic Acid (ETYA)	1780	-7.0	Forms two stable hydrogen bonds with Arginine at position 45 (ARG-45).
UBE2I	P63279	Eicosatetraynoic Acid (ETYA)	1780	-5.6	Forms one stable hydrogen bond with Arginine at position 62 (ARG-62).

6. Conclusion

The synthesis of high-resolution neuroimaging features with whole-blood transcriptomic biomarkers offers an advanced approach to managing acute ischemic stroke. By bridging macroscopic magnetic resonance imaging findings with microscopic metabolic gene signatures, this multi-scale framework can capture both the structural extent of brain

damage and the systemic physiological response, facilitating early diagnosis, precise therapeutic targeting, and robust risk stratification.

Despite this potential, several challenges hinder immediate clinical translation. A major barrier is the temporal and spatial discrepancy across modalities. Whole-blood transcriptomics reflects systemic cellular dynamics, which may not fully capture local molecular changes occurring within the ischemic brain parenchyma. Additionally, deep learning models can act as "black boxes," making it difficult to understand the complex, non-linear representations they learn. To address these limitations, incorporating Explainable AI (XAI) tools, such as Grad-CAM++ to highlight key imaging regions and SHAP value analysis to explain clinical feature contributions, can enhance model transparency. Validating these architectures across larger, multi-center cohorts is also essential to ensure generalizability across diverse scanners and patient demographics. Developing robust computational models that combine the speed of deep learning with the interpretability of classical statistical frameworks will support the creation of safe, reliable, and translationally viable decision-support systems for clinical stroke care.

References

1. Akyea, R. K., et al. (2022). Comparison of risk of serious cardiovascular events after hemorrhagic versus ischemic stroke: A population-based study. *Thrombosis and Haemostasis*, 122(11), 1921 - 1931.
2. Ben Alaya, I., Limam, H., & Kraiem, T. (2023). Automatic triaging of acute ischemic stroke patients for reperfusion therapies using artificial intelligence methods and multiple MRI features: A review. *Clinical Imaging*, 104, 109992.
3. Bian, X., Fu, H., Sun, W., Wang, N., Liu, H., & Han, W. (2025). Identification of post-translational modification-related biomarkers in ischemic stroke using bioinformatics and machine learning. *Scientific Reports*, 15, 39307. <https://doi.org/10.1038/s41598-025-22997-x>
4. Blank-Stein, N., & Mass, E. (2023). Macrophage and monocyte subsets in response to ischemic stroke. *European Journal of Immunology*, 53(10), e2250233.
5. Bondy-Chorney, E., et al. (2019). Nonhistone targets of KAT2A and KAT2B implicated in cancer biology (1). *Biochemistry and Cell Biology*, 97(1), 30 - 45.
6. Cao, Y., et al. (2020). LncRNA SNHG3 promotes autophagy-induced neuronal cell apoptosis by acting as a CeRNA for miR-485 to upregulate ATG7 expression. *Metabolic Brain Disease*, 35(8), 1361 - 1369.

7. Chen, R., et al. (2017). O-GlcNAc glycosylation of nNOS promotes neuronal apoptosis following glutamate excitotoxicity. *Cellular and Molecular Neurobiology*, 37(8), 1465 - 1475.
8. Chen, W., et al. (2021). Nanomedicines, an emerging therapeutic regimen for treatment of ischemic cerebral stroke: A review. *Journal of Controlled Release*, 340, 342–360.
9. Delong, J. H., et al. (2022). Inflammatory responses after ischemic stroke. *Seminars in Immunopathology*, 44(5), 625 - 648.
10. Thomalla, G., et al. (2018). MRI-guided thrombolysis for stroke with unknown time of onset. *New England Journal of Medicine*, 379(7), 611 - 622.
11. Lansberg, M. G., et al. (2019). Guidelines for the early management of acute ischemic stroke. *Stroke*, 50(12), e344 - e418.
12. Kamnitsas, K., et al. (2017). Efficient multi-scale 3D CNN with CRF for brain lesion segmentation. *Medical Image Analysis*, 36, 61 - 78.
13. Isensee, F., et al. (2021). nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18, 203 - 211.
14. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. *MICCAI*, 234 - 241.
15. Dvorianchikova, G., et al. (2009). Inactivation of astroglial NF-kappa B promotes survival of retinal neurons following ischemic injury. *European Journal of Neuroscience*, 30(2), 175 - 185.
16. Goncalves, A., et al. (2022). Thrombolytic tPA-induced hemorrhagic transformation of ischemic stroke is mediated by PKC β phosphorylation of occludin. *Blood*, 140(4), 388 - 400.
17. Götz, K. H., et al. (2019). Formation of the alarmones diadenosine triphosphate and tetraphosphate by ubiquitin- and ubiquitin-like activating enzymes. *Cellular Chemical Biology*, 26(11), 1535 - 1543.
18. Guan, I., et al. (2021). Synthetic thiol and selenol derived amino acids for expanding the scope of chemical protein synthesis. *Frontiers in Chemistry*, 9, 826764.
19. Gupta, M. K., et al. (2016). UBC9-mediated sumoylation favorably impacts cardiac function in compromised hearts. *Circulation Research*, 118(12), 1894 - 1905.