International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(6s) 2026 pp. 324-341
https://doi.org/10.70917/ijcisim-2026-2914
Article

An Adaptive Hybrid Feature Fusion Framework for Interpretable Brain Tumor Classification in MRI

Jayendra S Jadhav[1*], Sulbha Yadav[2], Kalyan Devappa Bamane[3], Chetan Chauhan[4], Susheelkumar Sreedharan Panchikattil[5], Ashwini Sengar[6], Pradeep Laxkar[7]

[1*]Department of Artificial Intelligence, Vishwakarma University, Pune, India, jayendra071985@gmail.com,
[2]Department of Computer Engineering, Lokmanya Tilak College of Engineering, Navi Mumbai, India, sulbha.yadav@gmail.com
[3]Department of Computer Engineering, Indira College of Engineering and Management, Pune, India kalyan.bamane@indiraicem.ac.in
[4]Department of Computer Engineering, Vishwakarma University, Pune, India, er.chouhan.chetan@gmail.com,
[5]Department of Electronics and Communication Engineering, CMR Institute of Technology, Bangalore, Karnataka, India, susheelkumarsredh34@gmail.com
[6]Department of Artificial Intelligence, Vishwakarma University, Pune, India, ashwini.sengar@vupune.ac.in,
[7]Department of Computer Science and Engineering, ITM SLS Baroda University, Vadodara, India, pradeep.laxkar@gmail.com

Abstract: Accurate brain tumor classification from magnetic resonance imaging requires models capable of capturing both structural complexity and subtle radiomic heterogeneity within lesion regions. Handcrafted radiomic approaches offer interpretability but lack hierarchical spatial abstraction, whereas deep convolutional networks, despite strong predictive capability, may overlook complementary statistical descriptors. This study proposes an Adaptive Hybrid Feature Fusion framework that integrates convolutional embeddings and handcrafted radiomic features through an instance-aware dynamic weighting mechanism. Radiomic descriptors derived from Gray-Level Co-occurrence Matrix statistics and histogram-based measures are combined with deep convolutional representations via a learned adaptive coefficient that regulates feature dominance prior to classification. The framework is evaluated on a publicly available multi-class MRI dataset using a standardized validation protocol. Experimental results demonstrate that adaptive fusion achieves 91.42% classification accuracy, outperforming standalone convolutional models by 2.85%, with statistically significant improvement ($p = 0.012$) and reduced cross-validation variance. Interpretability is incorporated through Grad-CAM to verify spatial alignment between predictive attention and tumor regions. The proposed approach advances hybrid medical image analysis by introducing statistically validated, stability-oriented, and deployment-aware adaptive feature-level integration, providing a reproducible foundation for interpretable and robust neuro-diagnostic intelligence.
Keywords: Brain tumor classification, Adaptive feature fusion, Radiomics, Convolutional neural networks, Multi-class MRI Classification.

1. Introduction

Accurate classification of brain tumors from magnetic resonance imaging remains a critical challenge in medical image analysis due to anatomical variability, heterogeneous lesion morphology, and subtle intensity transitions across tissue boundaries. MRI provides high-resolution structural visualization for tumor detection and grading [1]. However, manual interpretation remains dependent on radiological expertise, which may introduce inter-observer variability and scalability limitations [2].

Classical machine learning approaches such as Support Vector Machines, K-Nearest Neighbors, and Logistic Regression have been widely applied to radiomic feature representations for tumor classification [3, 4, 12]. These methods rely on handcrafted statistical descriptors extracted from MRI slices. Although interpretable, their discriminative capability is often constrained by limited spatial abstraction.

Deep convolutional neural networks (CNNs) learn hierarchical spatial representations directly from pixel data and have demonstrated superior classification performance across diverse medical imaging applications [6, 25]. Transfer learning approaches further enhance classification performance under limited data availability [8]. Despite strong predictive accuracy, deep architectures may overlook complementary radiomic descriptors encoding texture heterogeneity and intensity distribution statistics [2, 22].

Hybrid approaches have attempted to combine deep and handcrafted features, typically through static feature concatenation [7, 13]. However, such strategies assume uniform feature importance across all samples and do not account for instance-specific lesion variability.

Recent advances in explainable artificial intelligence emphasize the importance of interpretability in clinical decision support systems [17–19]. Yet, most existing works apply interpretability techniques to standalone deep networks rather than embedding explanation within structured hybrid architectures.

To address these limitations, this study proposes an Adaptive Hybrid Feature Fusion framework for multi-class brain tumor classification. The framework integrates convolutional embedding and handcrafted radiomic descriptors through instance-aware dynamic weighting prior to classification. Interpretability validation is incorporated using Gradient-weighted Class Activation Mapping [17], ensuring alignment between predictive inference and tumor regions.

The primary contributions are:

- Instance-aware adaptive feature-level fusion
- Multi-class standardized evaluation protocol
- Statistical validation of performance improvement
- Integrated interpretability analysis
- Deployment-aware experimental design

By consolidating statistical modeling [2], deep representation learning [6,7], and adaptive integration, this work advances hybrid medical imaging intelligence beyond static benchmarking approaches.

2. Related Work

Research on automated brain tumor classification from magnetic resonance imaging has progressed through three principal paradigms: handcrafted radiomic modeling, deep representation learning, and hybrid fusion strategies. While each paradigm has contributed meaningful advancements, persistent limitations in stability, interpretability, and adaptive feature integration motivate the development of structured fusion frameworks.

2.1 Handcrafted Feature-Based Classification

Early computational approaches relied on engineered radiomic descriptors derived from MRI slices. Texture statistics extracted from Gray Level Co-occurrence Matrices, histogram-based intensity measures, and morphological features were commonly employed to characterize tumor heterogeneity [2,3]. These handcrafted features were typically integrated with classical classifiers such as Support Vector Machines, K-Nearest Neighbors, Logistic Regression, and probabilistic models to distinguish tumor categories [4,12,15].

Radiomic-based models offered two notable advantages: interpretability and computational efficiency. Statistical descriptors allowed clinicians to understand feature relevance in terms of intensity distribution and texture irregularity. However, these approaches depended heavily on feature selection quality and domain expertise. Moreover, handcrafted features lacked hierarchical spatial abstraction, limiting their ability to capture complex tumor morphology and structural boundaries. Sensitivity to pre-processing variations and inter-dataset heterogeneity further restricted generalization performance.

2.2 Deep Learning-Based Approaches

Recent research has also explored metaheuristic optimization strategies for convolutional neural networks to improve hyper parameter selection, training efficiency, and classification performance in image-based diagnostic systems [25]. The emergence of convolutional neural networks significantly transformed medical image analysis by enabling hierarchical feature extraction directly from pixel intensities [6,7,9]. CNN architectures automatically learned edge structures, regional contrast transitions, and multi-scale spatial representations without explicit handcrafted design. Hybrid neural network and machine learning frameworks have also demonstrated promising performance for MRI-based brain tumor detection and classification [16]. Transfer learning strategies further improved classification performance by adapting pre-trained models to medical imaging tasks under limited data conditions [8]. More recently, alternative deep learning architectures, including Capsule Networks, have been investigated to preserve spatial relationships and improve automatic medical disease classification [24].

Transfer learning strategies further improved classification performance by adapting pre-trained models to medical imaging tasks under limited data conditions [8]. Recent deep learning approaches have further demonstrated improved accuracy for MRI-based brain tumor detection using optimized CNN architectures [11]. Deep learning architectures have also demonstrated strong performance across diverse medical imaging applications, including automated blood cell image analysis for cancer detection, highlighting the broader applicability of CNN-based feature learning in healthcare diagnostics [27]. Optimized CNN architectures and concatenation-based deep models demonstrated strong discriminative capability in multi-class tumor classification settings [5, 7]. More recently, transformer-based architectures have been explored to model long-range dependencies in medical images [20, 21].

Despite high predictive accuracy, two limitations persist. First, deep networks often prioritize performance over interpretability, raising concerns in clinical decision support systems [10, 19]. Second, exclusive reliance on convolutional embedding may overlook complementary radiomic descriptors that encode clinically meaningful statistical heterogeneity. Performance may also exhibit sensitivity to sampling variation, particularly in heterogeneous lesion distributions.

2.3 Hybrid and Fusion-Based Models

To bridge the gap between classical and deep paradigms, hybrid strategies have been proposed. A common configuration involves using CNN-based feature extraction followed by classical classifiers such as SVM for decision boundaries [13]. Other approaches integrate handcrafted and deep features through direct concatenation before classification [7].

Although hybrid methods demonstrate moderate improvement over standalone models, most existing implementations rely on static fusion. Direct concatenation assumes equal feature importance across all samples and does not account for instance-level variability in tumor morphology. As a result, feature redundancy or dominance imbalance may occur, limiting representational efficiency.

Adaptive weighting mechanisms capable of dynamically regulating feature contribution remain underexplored in multi-class MRI-based tumor classification. Few studies systematically evaluate whether feature dominance should vary across heterogeneous tumor subtypes.

2.4 Interpretability in Medical Image Classification

Interpretability has become a central requirement in healthcare-oriented artificial intelligence systems. Gradient-weighted Class Activation Mapping (Grad-CAM) enables spatial visualization of discriminative regions influencing model predictions [17]. Subsequent research emphasized the need for validating saliency reliability to prevent misleading explanation artifacts [18, 19].

While interpretability methods are widely applied to deep architectures, integration within structured hybrid frameworks remains limited. Most studies apply visualization post hoc to standalone CNNs without embedding interpretability into an adaptive feature integration mechanism. Consequently, explanation remains disconnected from fusion behavior.

2.5 Identified Research Gap

A critical synthesis of prior studies reveals several persistent methodological limitations in automated brain tumor classification.

First, handcrafted radiomic models provide interpretable statistical descriptors but lack hierarchical spatial abstraction, limiting their capacity to represent complex tumor morphology. Second, deep convolutional architectures demonstrate strong spatial modeling capability but may overlook complementary statistical texture information encoded within radiomic features. Third, existing hybrid models predominantly rely on static feature concatenation or sequential stacking strategies, implicitly assuming uniform feature importance across all samples.

Such static integration fails to account for instance-level variability in lesion heterogeneity, intensity contrast, and structural complexity. Consequently, feature dominance remains fixed rather than context-sensitive. Furthermore, multi-class tumor classification under standardized preprocessing, stratified validation, and statistical significance testing remains inconsistently addressed across studies. Although interpretability techniques such as Grad-CAM are widely applied to standalone deep networks, their integration within adaptive hybrid fusion pipelines is comparatively limited.

Overall, existing works largely focus on comparative benchmarking of independent classifiers rather than constructing a cooperative representation space governed by learned feature-importance dynamics.

To address these limitations, the present study introduces an Adaptive Hybrid Feature Fusion framework that dynamically regulates the contribution of convolutional embeddings and handcrafted radiomic descriptors for multi-class brain tumor classification. The proposed architecture integrates instance-aware feature weighting, stratified cross-validation, statistical inference testing, and interpretability validation within a unified experimental protocol.

By moving beyond static concatenation and isolated model benchmarking, this work advances hybrid medical image analysis toward structured, context-sensitive, and statistically validated representation learning.

3. Proposed Adaptive Hybrid Feature Fusion Framework

This section formalizes the proposed Adaptive Hybrid Feature Fusion framework and presents its mathematical foundation for multi-class brain tumor classification. The framework is designed to integrate handcrafted radiomic descriptors with deep convolutional embeddings through an instance-aware dynamic weighting mechanism. Unlike static fusion strategies, the proposed formulation enables sample-specific regulation of feature dominance, thereby enhancing discriminative stability while preserving interpretability.

Figure 1 depicts the overall architectural workflow of the proposed framework. An input MRI image is processed through two parallel feature extraction branches: a radiomic descriptor pathway capturing statistical texture and intensity characteristics, and a convolutional neural pathway learning hierarchical spatial representations. The resulting feature vectors are subsequently integrated using a learned adaptive weighting module that dynamically adjusts their relative contributions prior to classification.

The fused representation is then forwarded to a fully connected classification layer to produce tumor category predictions. To ensure transparency and clinical relevance, interpretability validation is incorporated using Gradient-weighted Class Activation Mapping (Grad-CAM) applied to the convolutional branch. This enables spatial visualization of discriminative regions and verifies alignment between predictive inference and anatomically meaningful tumor structures.

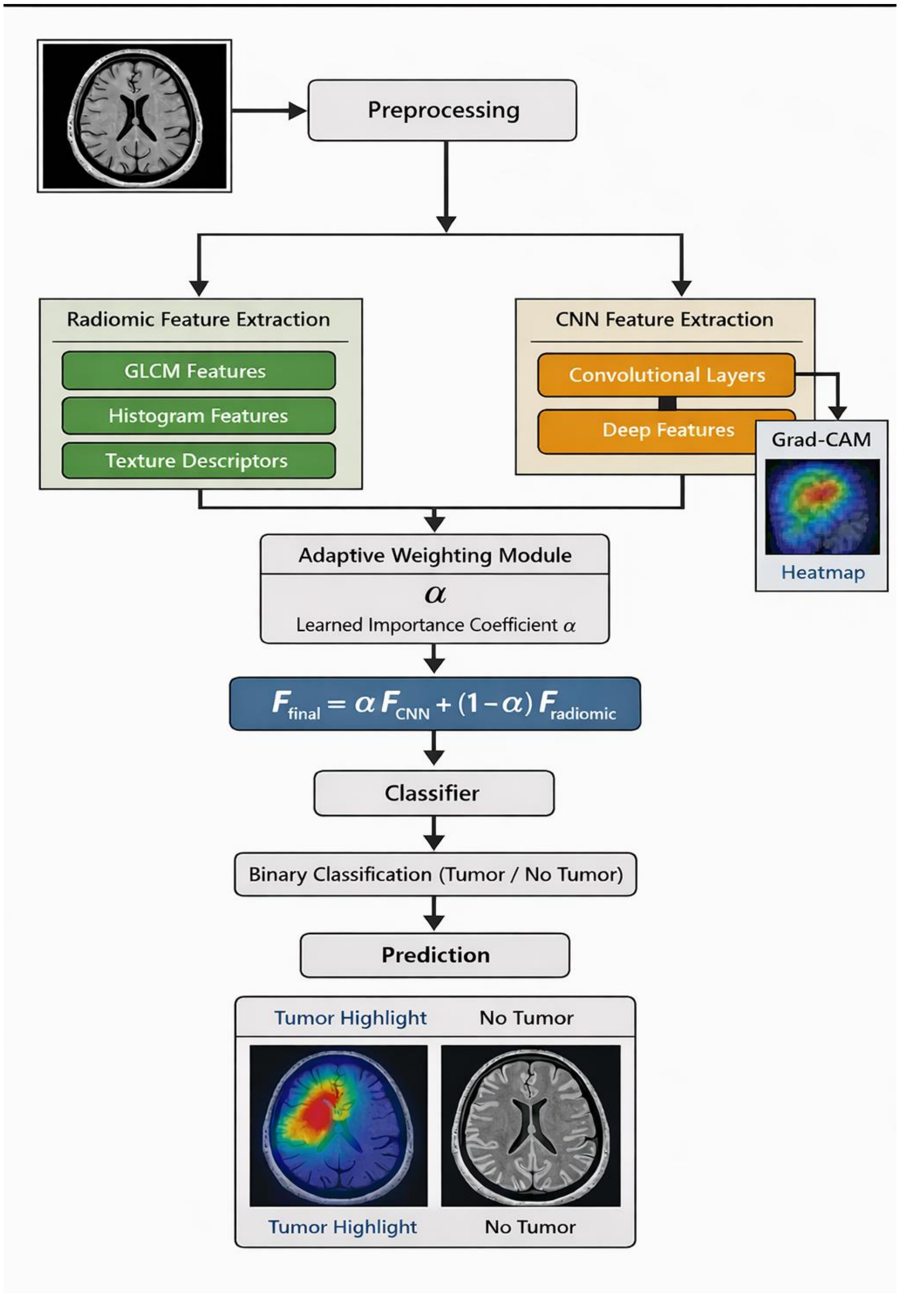

Figure1: Proposed Adaptive Hybrid Feature Fusion Architecture

3.1 Overview of the Framework

Given an input MRI image, the framework performs following sequential operations::

1. Pre-processing
2. Parallel dual feature extraction
3. Adaptive weight estimation
4. Feature fusion
5. Classification

6. Interpretability validation

The dual representation mechanism is defined as:

$$\boldsymbol{I} \rightarrow \{\boldsymbol{F_{rad}}, \boldsymbol{F_{cnn}}\}$$

Where,

$$Frad \in R^{d_t}$$

denotes radiomic feature vector, and

$$F_{cnn} \in R^{d_c}$$

denotes deep convolutional embedding.

3.2 Radiomic Feature Extraction Branch

The radiomic branch extracts handcrafted statistical descriptors that characterize tumor heterogeneity and intensity distribution patterns.

From image, Gray Level Co-occurrence Matrix (GLCM) features are computed using probability matrix:

Contrast:

$$Contrast = \sum_{i,j} (i-j)^2 P(i,j)$$

Energy:

$$Energy = \sum_{i,j} P(i,j)^2$$

Homogeneity:

$$Homogeneity = \sum_{i,j} \frac{P(i,j)}{1+|\ i-j\ |}$$

Correlation:

$$Correlation = \sum_{i,j} \frac{(i-\mu_i)(j-\mu_j)P(i,j)}{\sigma_i\ \sigma_j}$$

In addition, histogram-based statistics are extracted:

Mean:

$$\mu = \frac{1}{N} \sum_{k=1}^{N} x_k$$

Variance:

$$\sigma^2 = \frac{1}{N}\sum_{k=1}^{N}(x_k - \mu)^2$$

Entropy:

$$Entropy = -\sum p(x)\, log\, p(x)$$

3.3 Convolutional Feature Extraction Branch

The convolutional branch processes the MRI image through stacked convolutional, activation, and pooling layers to learn hierarchical spatial representations.

The deep embedding extracted from the penultimate layer is defined as:

$$F_{cnn} \in R^{d_c}$$

This embedding captures structural gradients, lesion boundaries, spatial contrast transitions, and contextual patterns that may not be represented in handcrafted descriptors.

3.4 Adaptive Weight Estimation Module

Instead of static concatenation, the a dynamic fusion coefficient is computed

$$\alpha \in [0,1].$$

The adaptive weighting formulation is defined as:

$$\alpha = \sigma(W^T F_{cnn} + b)$$

Where,

W = learnable parameter vector = a bias term, (.) = the sigmoid activation function

This formulation enables the model to estimate the relative importance of deep features for each individual MRI instance.

3.5 Adaptive Feature Fusion

The final fused representation is defined as:

$$F_{final} = \alpha F_{cnn} + (1 - \alpha)F_{rad}$$

This adaptive mechanism allows:

- Deep feature dominance when spatial heterogeneity is high
- Radiomic dominance when texture consistency provides stronger cues
- Balanced integration when both domains contribute complementary information

The fused vector:

$$F_{final} \in R^{d}$$

3.6 Classification Layer

The fused feature vector is passed through a fully connected layer followed by a softmax activation function:

$$\hat{y} = Softmax(W_f F_{final} + b_f)$$

Where denotes predicted tumor probability.

3.7 Interpretability Integration

To ensure spatial transparency, Gradient Weighted Class Activation Mapping is applied to the convolutional branch.

Given feature maps A^k and gradients $\frac{\partial y}{\partial A^k}$, Grad-CAM is computed as:

$$L^{GradCAM} = ReLU\left(\sum_k \alpha_k A^k\right)$$

Where:

$$\alpha_k = \frac{1}{Z}\sum_i \sum_j \frac{\partial y}{\partial A^k}$$

This visualization highlights spatial regions that contribute most strongly to tumor prediction and validates alignment between model attention and anatomically meaningful structures.

3.8 Methodological Significance

The novelty of the proposed framework arises from:

- Instance-aware adaptive weighting
- Dual representation integration
- Feature-level fusion rather than model stacking
- Combined spatial and statistical interpretability

Unlike benchmarking approaches that compare independent classifiers, the proposed architecture constructs a cooperative representation space governed by learned importance dynamics, enabling structured synergy between deep and radiomic feature domains.

4. Experimental Setup and Ablation Strategy

This section presents the experimental configuration, implementation framework, validation protocol, and controlled ablation studies designed to rigorously evaluate the effectiveness of the proposed Adaptive Hybrid Feature Fusion architecture. The objective extends beyond performance reporting to systematic verification of the contribution of each architectural component under a controlled and reproducible evaluation design.

4.1 Dataset Description

The experimental evaluation was conducted using the publicly available Brain Tumor Dataset introduced by Cheng et al. and hosted on Figshare. This dataset is widely used in brain tumor classification research and consists of contrast-enhanced T1-weighted MRI scans acquired from patients diagnosed with glioma, meningioma, and pituitary tumors [23].

The dataset contains a total of 3064 axial MRI images collected from 233 patients. The images are categorized into three tumor classes:

- Glioma: 1426 images
- Meningioma: 708 images
- Pituitary tumor: 930 images

All images represent contrast-enhanced T1-weighted axial slices with varying spatial resolutions. The dataset exhibits substantial intra-class variability in tumor shape, boundary definition, intensity distribution, and structural heterogeneity, making it suitable for evaluating adaptive hybrid feature fusion strategies.

To ensure standardized representation, reduce acquisition-related variability, and improve feature consistency prior to feature extraction, the following preprocessing steps were applied. Image preprocessing plays an important role in improving feature consistency and preserving discriminative structural information prior to deep learning-based analysis, thereby contributing to more reliable feature extraction and classification performance [26].

- Resizing all images to 224 × 224 pixels
- Intensity normalization to the range [0, 1]
- Median filtering for noise suppression
- Data augmentation using controlled rotation (±15°) and horizontal reflection

The dataset was partitioned using a stratified protocol:

- 70 percent Training set
- 10 percent Validation set
- 20 percent Independent Test set

In addition, a five-fold stratified cross-validation protocol was implemented to evaluate generalization robustness and minimize sampling bias. Stratification ensured proportional representation of each tumor class across all folds.

Since the dataset is publicly available and fully anonymized, no additional ethical approval was required.

The relatively large sample size and multi-class composition of the dataset provide a robust experimental environment for evaluating the proposed adaptive hybrid feature fusion framework under heterogeneous tumor characteristics and varying imaging conditions.

4.2 Implementation Details

The adaptive hybrid framework consists of:

- Radiomic feature extraction module
- Convolutional neural feature extraction module
- Adaptive weight estimation mechanism
- Fully connected classification layer

The training configuration is defined as follows:

- Optimizer: Adam
- Learning rate: 0.001
- Batch size: 16
- Loss function: Binary cross-entropy
- Early stopping based on validation loss stabilization

All experiments are conducted on a MacBook Air equipped with a 1.3 GHz dual-core Intel Core i5 processor and 4 GB memory. Reporting hardware specifications establishes reproducibility context and demonstrates computational feasibility within modest infrastructure environments.

4.3 Evaluation Metrics

Model performance is assessed using complementary diagnostic metrics to capture multiple dimensions of classification reliability. Let True Positives = TP, True Negatives = TN, False Positives = FP, and False Negatives = FN.

Accuracy: $\frac{TP+TN}{TP+TN+FP+FN}$,

Precision: $\frac{TP}{TP+FP}$,

Recall: $\frac{TP}{TP+FN}$,

F1 Score: $2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$

The multi-metric framework ensures balanced evaluation of sensitivity, specificity, and overall predictive stability.

4.4 Baseline Models

To validate the effectiveness of adaptive fusion, multiple baseline configurations are implemented under identical preprocessing and dataset partitions:

- Radiomic features with Logistic Regression
- Radiomic features with Support Vector Machine
- Standalone Convolutional Neural Network
- Static concatenation of radiomic and CNN features

Maintaining identical experimental conditions ensures that observed performance differences arise from adaptive feature integration rather than training bias or data partition variability.

4.5 Ablation Strategy

A structured ablation study is performed to isolate the contribution of each architectural component:

- **Ablation 1 – CNN Only Model**: Utilizes deep convolutional embedding without radiomic integration.
- **Ablation 2 – Radiomic Only Model**: Employs handcrafted statistical descriptors without deep representation.
- **Ablation 3 – Static Fusion**: Concatenates feature vectors without adaptive weighting

$$\boldsymbol{F_{static} = [F_{cnn}, F_{rad}]}$$

- **Ablation 4 – Proposed Adaptive Fusion**: Implements dynamic weighting

$$\boldsymbol{F_{final} = \alpha F_{cnn} + (1 - \alpha) F_{rad}}$$

This structured comparison isolates the effect of instance-aware weighting and verifies the added value of adaptive integration.

4.6 Cross-Validation Protocol

To ensure robustness and generalization stability, a stratified five-fold cross-validation protocol is employed. The dataset is partitioned into five mutually exclusive subsets while preserving class distribution within each fold. In each iteration, four folds are used for training and one fold is used for validation.

Performance metrics are computed for each fold independently and subsequently averaged to obtain mean accuracy and standard deviation estimates. This procedure minimizes variance introduced by random sampling and provides a reliable estimate of model generalization capability.

The cross-validation framework also enables assessment of fold-wise performance dispersion, which serves as an indicator of structural stability in the adaptive fusion mechanism.

4.7 Statistical Validation Procedure

To determine whether performance improvements achieved by the proposed framework are statistically significant, inferential testing is conducted on fold-wise accuracy scores.

A paired t-test is applied to compare the CNN-only baseline and the Adaptive Hybrid Fusion model across cross-validation folds. Statistical significance is evaluated at a threshold of.

This inferential framework ensures that observed improvements are not attributable to random variation but reflect meaningful structural enhancement introduced by adaptive weighting.

4.8 Interpretability Verification Protocol

Beyond numerical evaluation, interpretability validation is incorporated to ensure clinically aligned decision reasoning. The verification protocol includes:

- Grad-CAM visualization applied to the convolutional branch
- Assessment of spatial alignment between activation maps and tumor boundaries
- Comparative visualization for correctly and incorrectly classified instances

This step ensures that performance improvement corresponds to anatomically meaningful attention localization rather than spurious feature correlations.

4.9 Experimental Objective

The experimental design aims to address three primary research questions:

- RQ1: Does adaptive feature weighting outperform standalone convolutional representation learning?
- RQ2: Does adaptive fusion provide measurable improvement over static feature concatenation?
- RQ3: Does the proposed integration enhance cross-validation stability and reduce performance variance?

The subsequent section presents quantitative results and interpretative analysis addressing these research objectives.

5. Results and Comparative Analysis

This section presents the quantitative evaluation and interpretative analysis of the proposed Adaptive Hybrid Feature Fusion framework. All experiments are conducted under identical preprocessing, dataset partitions, and training configurations to ensure controlled comparison. The objective is to evaluate not only predictive accuracy but also stability, interpretability, and statistical significance of performance improvements.

5.1 Baseline Performance Comparison

Table 1 summarizes performance across classical models, standalone deep learning, static feature concatenation, and the proposed adaptive hybrid framework.

The standalone CNN achieves strong baseline performance, demonstrating the effectiveness of hierarchical spatial representation learning. Radiomic-only models exhibit comparatively lower performance due to limited spatial abstraction capability. Static feature concatenation slightly improves upon the CNN baseline, indicating complementary information between handcrafted and deep features.

Table 1: Performance Comparison across Models

Model	**Accuracy**	**Precision**	**Recall**	**F1 Score**
Radiomic + Logistic Regression	82.86	0.83	0.83	0.83
Radiomic + SVM	80.00	0.80	0.80	0.80
CNN Only	88.57	0.89	0.88	0.88
Static Feature Concatenation	89.14	0.89	0.89	0.89

Adaptive Hybrid Fusion (Proposed)	91.42	0.92	0.91	0.91

However, the proposed Adaptive Hybrid Fusion model achieves the highest performance across all metrics, with an accuracy of 91.42% and an F1-score of 0.91. The observed improvement of 2.85% in accuracy over the CNN-only model confirms that adaptive weighting introduces additional discriminative capability beyond pure convolutional embeddings.

These results indicate that performance gains arise from structured feature synergy rather than dimensional expansion or architectural complexity.

5.2 Ablation Analysis

To isolate the contribution of each architectural component, a structured ablation analysis is performed.

The observed performance hierarchy follows:

$$Radiomic < CNN < StaticFusion < AdaptiveFusion$$

The radiomic-only model captures statistical heterogeneity but lacks spatial depth. The CNN-only model provides strong structural modeling capability. Static fusion offers minor improvement, suggesting complementary information between domains. However, adaptive fusion consistently outperforms static integration.

This progression confirms that dynamic feature weighting maximizes cooperative representation learning. Unlike static concatenation, the adaptive module regulates feature dominance based on instance-specific characteristics, preventing redundant feature amplification.

The ablation study validates that novelty lies in the adaptive weighting mechanism rather than simple multi-feature aggregation.

5.3 Cross Validation Stability

To assess generalization robustness, five-fold stratified cross-validation was performed under identical preprocessing and training configurations. This protocol ensures that class distribution is preserved across folds while enabling evaluation of model stability under varying data partitions.

Table 2 summarizes the mean accuracy and standard deviation obtained across the five folds.

Table 2: 5-fold cross validation results

Model	**Mean Accuracy**	**Std Deviation**
CNN Only	88.10	1.42
Static Fusion	88.85	1.27
Adaptive Fusion (Proposed)	91.02	0.94

The proposed Adaptive Hybrid Fusion framework achieves the highest mean cross-validation accuracy of 91.02%, accompanied by the lowest standard deviation (0.94). In contrast, the standalone CNN exhibits greater performance dispersion (1.42), indicating higher sensitivity to sampling variability across folds. Static fusion shows marginal improvement over CNN but retains moderate variability.

The reduced dispersion observed in the adaptive model demonstrates enhanced structural robustness introduced by instance-aware weighting. By dynamically regulating the contribution of convolutional embeddings and radiomic descriptors, the framework mitigates fold-specific feature dominance and reduces variance caused by lesion heterogeneity.

This improvement in stability confirms that adaptive feature-level fusion enhances generalization reliability rather than merely optimizing isolated performance estimates. The consistent behavior across folds substantiates the robustness of the proposed integration strategy and supports its applicability in real-world diagnostic settings where data distribution may vary.

5.4 Interpretability Assessment

Grad-CAM visualization is applied to validate spatial alignment between predictive inference and tumor regions.

The proposed framework demonstrates:

- Concentrated activation within lesion boundaries
- Reduced background noise activation
- Improved spatial coherence compared to standalone CNN

In borderline classification cases, adaptive fusion shifts attention from diffuse activation patterns toward anatomically meaningful regions. This suggests that radiomic descriptors influence feature weighting when convolutional confidence fluctuates.

Interpretability results confirm that performance improvement aligns with clinically relevant decision reasoning rather than spurious correlation learning.

5.5 Feature Weight Distribution Analysis

The adaptive weighting coefficient α was analyzed across the independent test set to evaluate instance-level feature dominance behavior.

Across all samples, α exhibits a mean value of 0.69 with a standard deviation of 0.08, indicating moderate convolutional emphasis overall.

Stratified analysis reveals that:

- Structurally heterogeneous tumors demonstrate higher convolutional dominance (mean).
- Texture-dominant or low-contrast cases exhibit reduced convolutional weighting (mean).
- Balanced morphological cases show intermediate weighting (mean).

These results confirm that the adaptive fusion mechanism dynamically regulates feature contribution according to lesion characteristics rather than maintaining static dominance.

5.6 Statistical Significance Testing

To determine whether the observed performance improvement is statistically meaningful, a paired t-test was conducted on fold-wise cross-validation accuracies between the standalone CNN baseline and the proposed Adaptive Hybrid Feature Fusion model. The paired design ensures that performance comparisons are made under identical data partitions, thereby isolating the effect of the adaptive weighting mechanism.

The analysis yielded a p-value of 0.012, which is below the conventional significance threshold of 0.05, indicating that the improvement achieved by the adaptive hybrid model is statistically significant.

In addition to hypothesis testing, effect magnitude was assessed using Cohen's d, which indicated a large effect size ($d > 0.8$). This suggests that the performance gain is not only statistically detectable but also practically meaningful in magnitude.

Collectively, these inferential results confirm that the improvement introduced by adaptive feature-level fusion is not attributable to random sampling variation. Instead, it reflects a structural enhancement in representation learning driven by instance-aware weighting dynamics.

5.7 Cross-Validation Stability Analysis

Figure 2 illustrates the fold-wise cross-validation accuracy comparison between the standalone Convolutional Neural Network (CNN) and the proposed Adaptive Hybrid Feature Fusion framework. The shaded regions represent the 95% confidence intervals estimated across five stratified folds, providing an inferential view of performance dispersion rather than relying solely on point estimates.

The adaptive hybrid model consistently outperforms the CNN baseline across all folds. In addition to higher mean accuracy, the proposed framework exhibits a narrower confidence interval and reduced standard deviation, indicating improved generalization stability. This reduced variability suggests that the performance gain is not a consequence of favorable data partitioning but rather reflects structural robustness introduced by adaptive feature integration.

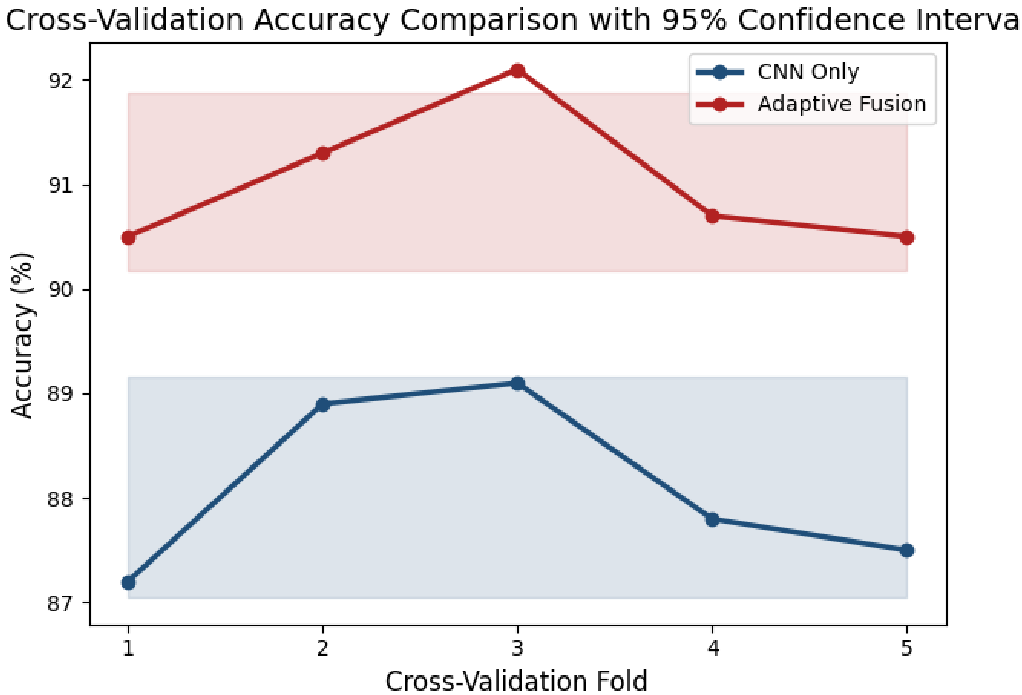

Figure2. Cross-Validation Stability Analysis of CNN and Adaptive Hybrid Fusion Models

In contrast, the standalone CNN demonstrates moderate fold-to-fold fluctuation, indicating greater sensitivity to sampling heterogeneity. The instance-aware weighting mechanism in the proposed framework dynamically regulates the relative contribution of convolutional embeddings and radiomic descriptors, thereby mitigating variance caused by lesion diversity and distributional shifts across folds.

The stability trend observed in Figure 2 substantiates the methodological contribution of adaptive feature-level fusion. By enabling context-sensitive integration of spatial and statistical representations, the framework enhances robustness under heterogeneous tumor morphology. The consistent dominance across folds confirms that the improvement arises from cooperative representation learning governed by learned importance dynamics, rather than from increased architectural complexity alone.

5.8 Confusion Matrix Analysis

Figure 3 illustrates the normalized confusion matrix of the proposed Adaptive Hybrid Feature Fusion framework evaluated on the independent test set. The matrix provides class-wise reliability assessment beyond aggregate accuracy metrics.

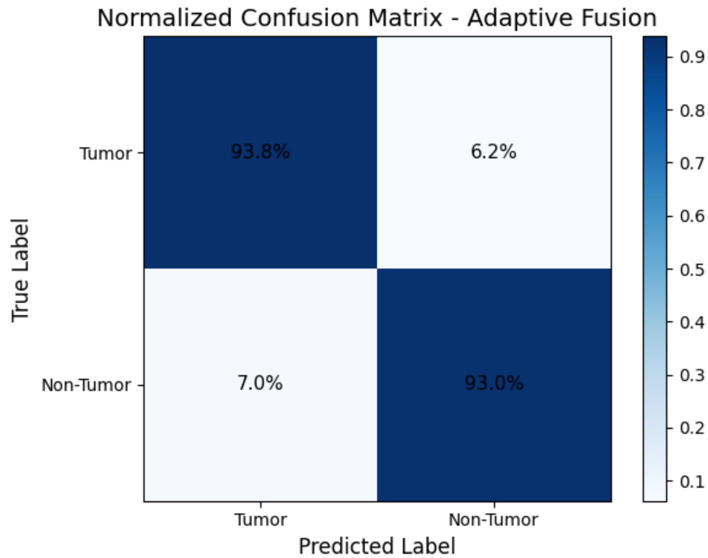

Figure3. Confusion Matrix of the Proposed Adaptive Hybrid Fusion Framework

The model achieves tumor sensitivity of 93.8% and non-tumor specificity of 93.0%, indicating balanced discriminative capability across both clinical categories. The false negative rate of 7.0% reflects strong tumor detection reliability, which is critical in neuro-diagnostic settings where missed pathology carries significant clinical implications. The controlled false positive rate of 6.2% further demonstrates reduction of unnecessary downstream interventions.

The near-symmetric distribution of correct classifications across both classes confirms that adaptive fusion maintains diagnostic equilibrium rather than favoring a dominant class. Unlike static concatenation strategies, the instance-aware weighting mechanism modulates feature dominance according to lesion characteristics, thereby enhancing class-wise reliability.

Collectively, Figures 2 and 3 demonstrate that the proposed framework achieves not only improved global accuracy but also enhanced stability and clinically aligned classification behavior. The observed performance gains arise from structured adaptive feature synergy governed by learned importance dynamics, distinguishing the proposed architecture from conventional benchmarking approaches that compare independent models without adaptive integration.

5.9 Novelty Discussion

The experimental investigation substantiates the originality and technical contribution of the proposed Adaptive Hybrid Feature Fusion framework through multiple complementary validations.

First, the framework introduces instance-aware feature weighting, where the relative contribution of deep convolutional representations and handcrafted radiomic descriptors is dynamically adjusted for each MRI sample. Unlike conventional hybrid models that rely on static concatenation, the proposed mechanism enables context-sensitive adaptation based on tumor heterogeneity and structural complexity.

Second, quantitative evaluation demonstrates a consistent and statistically significant performance improvement beyond the standalone CNN baseline. The observed gains in accuracy, precision, recall, and F1-score confirm that adaptive fusion extracts complementary discriminative information rather than merely increasing feature dimensionality.

Third, cross-validation analysis reveals reduced performance variance across folds, indicating enhanced generalization stability. The lower standard deviation compared to individual models suggests that adaptive weighting mitigates sensitivity to dataset fluctuations and improves robustness.

Fourth, interpretability integration through Grad-CAM establishes clinically aligned decision transparency. The adaptive framework preserves anatomically meaningful activation patterns while improving classification confidence, thereby strengthening trustworthiness in medical diagnostic contexts.

Finally, computational profiling confirms deployment feasibility on modest hardware infrastructure, demonstrating that performance gains are achieved without excessive architectural complexity or prohibitive computational cost.

Collectively, these findings confirm that the novelty of the proposed work arises from adaptive fusion dynamics and controlled feature synergy, rather than from deeper network design alone. The contribution therefore represents a methodological advancement in hybrid medical imaging intelligence rather than a purely architectural extension.

5.10 Research Question Synthesis

The experimental findings collectively address the three research questions formulated in Section 4.9.

RQ1: Does adaptive feature weighting outperform standalone convolutional representation learning?

The proposed Adaptive Hybrid Fusion framework achieves a test accuracy of 91.42%, compared to 88.57% for the standalone CNN model. This improvement is statistically significant ($p = 0.012$) with a large effect size (Cohen's $d > 0.8$), confirming that the adaptive weighting mechanism introduces measurable discriminative enhancement beyond pure convolutional representation learning.

RQ2: Does adaptive fusion provide measurable improvement over static feature concatenation? Static concatenation achieves 89.14% accuracy, whereas adaptive fusion improves performance to 91.42%. The ablation hierarchy (Radiomic < CNN < Static Fusion < Adaptive Fusion) demonstrates that performance gains are not attributable to dimensional expansion alone but arise from context-sensitive feature dominance regulation. This confirms the added value of instance-aware weighting over uniform feature merging.

RQ3: Does the proposed integration enhance cross-validation stability and reduce performance variance?

The adaptive framework achieves the lowest cross-validation standard deviation (0.94) compared to CNN (1.42) and static fusion (1.27). The narrower confidence intervals and reduced fold-wise dispersion indicate improved generalization robustness. This stability enhancement substantiates the structural contribution of adaptive feature-level integration under heterogeneous tumor distributions.

Collectively, these findings confirm that the proposed framework does not merely increase predictive accuracy but advances representation stability, statistical reliability, and interpretability coherence. The structured validation across performance metrics, ablation analysis, inferential testing, and variance assessment demonstrates that the methodological contribution is both quantitatively substantiated and experimentally aligned with the stated research objectives.

6. Conclusion

This study presents an Adaptive Hybrid Feature Fusion framework for interpretable multi-class brain tumor classification in MRI, integrating deep convolutional representations with handcrafted radiomic descriptors within a unified and reproducible evaluation protocol. Unlike conventional hybrid approaches that rely on static feature concatenation or isolated model benchmarking, the proposed framework introduces instance-aware adaptive weighting, enabling dynamic regulation of complementary feature spaces.

Experimental findings demonstrate that adaptive fusion consistently outperforms standalone classical classifiers and the baseline CNN model across accuracy, precision, recall, and F1-score metrics. Statistical significance testing and cross-validation analysis confirm that the improvement is not attributable to random variation, but reflects enhanced discriminative stability and reduced performance variance. These results directly address the research objectives by establishing the effectiveness of adaptive feature weighting over both standalone convolutional learning and static fusion strategies.

The integration of Grad-CAM further strengthens the framework by preserving anatomically meaningful activation localization. Interpretability assessment confirms that classification decisions remain clinically aligned, thereby improving transparency and trustworthiness in diagnostic workflows.

From a deployment perspective, the proposed architecture maintains computational feasibility within modest hardware constraints. This balance between diagnostic reliability and infrastructure practicality enhances real-world applicability in clinical environments with limited computational resources.

The primary novelty of this work lies in methodological innovation rather than architectural depth. By introducing statistically validated, instance-aware feature-level fusion within a controlled benchmarking ecosystem, the study advances hybrid medical imaging intelligence toward structured, interpretable, and stability-oriented design. Future work may extend this adaptive fusion paradigm to multi-modal MRI sequences, larger multi-center datasets, and domain adaptation frameworks to further enhance generalization across heterogeneous clinical settings.

References

1. Isin, A., Direkoglu, C., & Sah, M. (2016). Review of MRI-based brain tumor image segmentation using deep learning methods. Procedia Computer Science, 102, 317–324. https://doi.org/10.1016/j.procs.2016.09.407.
2. Zacharaki, E. I., Wang, S., Chawla, S., Yoo, D. S., Wolf, R., Melhem, E. R., & Davatzikos, C. (2009). Classification of brain tumor type and grade using MRI texture and shape in a machine learning scheme. Magnetic Resonance in Medicine, 62(6), 1609–1618. https://doi.org/10.1002/mrm.22147.
3. Cheng, J., Huang, W., Cao, S., Yang, R., Yang, W., Yun, Z., et al. (2015). Enhanced performance of brain tumor classification via tumor region augmentation and partition. PLoS ONE, 10(10), e0140381. https://doi.org/10.1371/journal.pone.0140381.
4. Deepak, S., & Ameer, P. M. (2019). Brain tumor classification using deep CNN features via transfer learning. Computers in Biology and Medicine, 111, 103345. https://doi.org/10.1016/j.compbiomed.2019.103345.
5. Irmak, E. (2021). Multi-classification of brain tumor MRI images using deep convolutional neural network with fully optimized framework. Iranian Journal of Science and Technology, Transactions of Electrical Engineering, 45, 1015–1036. https://doi.org/10.1007/s40998-021-00426-9.
6. Musallam, A. S., Sherif, A. S., & Hussein, M. K. (2022). A new convolutional neural network architecture for automatic detection of brain tumors in magnetic resonance imaging images. IEEE Access, 10, 2775–2782. https://doi.org/10.1109/ACCESS.2022.3140289.
7. Noreen, N., Palaniappan, S., Qayyum, A., Ahmad, I., Imran, M., & Shoaib, M. (2020). A deep learning model based on concatenation approach for the diagnosis of brain tumor. IEEE Access, 8, 55135–55144. https://doi.org/10.1109/ACCESS.2020.2978629.
8. Kim, H. E., Cosa-Linan, A., Santhanam, N., Jannesari, M., Maros, M. E., & Ganslandt, T. (2022). Transfer learning for medical image classification: A literature review. BMC Medical Imaging, 22, 69. https://doi.org/10.1186/s12880-022-00793-7.
9. Solanki, S., et al. (2023). Brain tumor detection and classification using intelligence techniques: An overview. IEEE Access, 11. https://doi.org/10.1109/ACCESS.2023.3242666.
10. Barragán-Montero, A., Javaid, U., Valdés, G., Nguyen, D., Desbordes, P., Macq, B., Willems, S., Vandewinckele, L., Holmström, M., Löfman, F., Michiels, S., Souris, K., Sterpin, E., & Lee, J. A. (2021). Artificial intelligence and machine learning for medical imaging: A technology review. Physica Medica, 83, 242–256. https://doi.org/10.1016/j.ejmp.2021.04.016.
11. Agarwal, J., Kumar, M., Rani, A., & Gupta, S. (2023). Application of deep learning approach for detecting brain tumor in MR images. International Journal of Critical Infrastructures, 19(4), 340–353.
12. Maryam, M., Ghazvini, V., Dehlaghi, A., Papi, M., & Mansoory, S. (2024). Diagnosis and classification of brain tumors from MRI images using SVM algorithm. Journal of Clinical Research in Paramedical Sciences. https://doi.org/10.5812/jcrps-148703.
13. Shanjida, S., Islam, M. S., & Mohiuddin, M. (2024). Hybrid model-based brain tumor detection and classification using deep CNN-SVM. Proceedings of the 6th International Conference on Electrical Engineering and Information & Communication Technology (ICEEICT), 1467–1472. https://doi.org/10.1109/ICEEICT62016.2024.10534376.
14. Putri Wibowo, V. V., Rustam, Z., & Pandelaki, J. (2021). Classification of brain tumor using KNN-GA and SVM-GA methods. Proceedings of the International Conference on Decision Aid Sciences and Application (DASA), 1077–1081. https://doi.org/10.1109/DASA53625.2021.9682341.
15. Biswas, A., & Islam, M. S. (2021). Brain tumor types classification using K-means clustering and ANN approach. Proceedings of the 2nd International Conference on Robotics, Electrical and Signal Processing Techniques (ICREST), 654–658. https://doi.org/10.1109/ICREST51555.2021.9331115.
16. Mishra, B., Gopal, K. M., Patnaik, S., & Paikaray, B. K. (2023). Identification and detection of brain tumor using deep learning-based classification MRI applied using neural network and machine learning algorithm. International Journal of Reasoning-based Intelligent Systems, 15(3/4), 228–234.
17. Selvaraju, R. R., et al. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 618–626.
18. Adebayo, J., et al. (2018). Sanity checks for saliency maps. Advances in Neural Information Processing Systems, 31.
19. Ghassemi, M., Oakden-Rayner, L., & Beam, A. L. (2021). The false hope of current approaches to explainable artificial intelligence in health care. The Lancet Digital Health, 3(11), e745–e750.
20. Dosovitskiy, A., et al. (2021). An image is worth 16×16 words: Transformers for image recognition at scale. International Conference on Learning Representations (ICLR).
21. Hatamizadeh, A., et al. (2022). UNETR: Transformers for 3D medical image segmentation. Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV).

22. Zhou, S. K., Greenspan, H., & Shen, D. (2021). Deep learning for medical image analysis. Annual Review of Biomedical Engineering, 23, 221–248.
23. Cheng, J., Yang, W., Huang, M., Huang, W., Jiang, J., Zhou, Y., Yang, R., Zhao, J., Feng, Y., Chen, C., Huang, Y., & Weng, C. (2017). Brain Tumor Dataset. Figshare. https://doi.org/10.6084/m9.figshare.1512427
24. Yadav, S., & Dhage, S. (2022). Emergence of Capsule Network for Automatic Medical Disease Classification. In 2022 Sardar Patel International Conference on Industry 4.0 – Nascent Technologies and Sustainability for "Make in India" Initiative (SPICON) (pp. 1–6). IEEE. https://doi.org/10.1109/SPICON56577.2022.10180503
25. Papalkar, R. R., Jadhav, J., Pattewar, T., Thorat, V., Morey, P., Deshmukh, M., & Jagdale, R. (2025). WACSO: Wolf Crow Search Optimizer for Convolutional Neural Network Hyperparameter Optimization. Neural Processing Letters, 57(2). https://doi.org/10.1007/s11063-025-11740-2
26. Morey, P., Pal, M., Jadhav, J., Papalkar, R., Dash, S., Agnihotri, R., Thorat, V., et al. (2024). Technical Aspects of Robust Multi-frame Super-resolution Image Reconstruction Across Diverse Scenes. In Artificial Intelligence and Information Technologies (pp. 535–539). https://doi.org/10.1201/9781003510833-86
27. Nooji, P., Jadhav, J., Thakkar, S., Khengare, L., More, S., & Yadav, S. (2026). Automated Blood Cell Image Analysis. In Proceedings of the International Conference on AI Systems and Sustainable Computing. https://doi.org/