

A Deep Learning Framework for Brain Tumor Classification Using VGG16-Based Autoencoder with BiLSTM Feature Refinement

S. Jansi¹, P. T. Bharathi²,
A. B. Feroz Khan^{3,*} and R. Jayanthi⁴

¹Department of Computer Science and Technology, Madanapalle Institute of Technology & Science, Madanapalle – 517325, India

²Department of Computer Science and Engineering, Siddaganga Institute of Technology, Tumakuru, India

³Department of Computer Science, Syed Hameedha Arts and Science College, Kilakarai, Tamil Nadu, India

⁴Department of MCA, Dayananda Sagar College of Engineering, Kumaraswamy Layout, Bangalore 560078, India

*E-mail: abferozkhan@gmail.com

Received for publication
May 23, 2025

Abstract

Brain tumor (BT) detection from magnetic resonance imaging (MRI) images plays a critical role in early diagnosis and effective treatment planning. This research proposes a novel deep learning framework that integrates a VGG16-based Convolutional Autoencoder (CAE) with a Bi-Directional Long Short-Term Memory (BiLSTM) network for accurate and robust BT classification. The VGG16-CAE is employed for hierarchical spatial feature extraction, capturing both low-level and high-level representations of MRI images, while the BiLSTM module models long-range dependencies within the extracted features, enhancing classification performance. The model was evaluated using the publicly available Kaggle Brain MRI Dataset with systematic hyperparameter tuning and fivefold cross-validation to ensure reproducibility and generalization. The proposed framework achieved superior performance, with an average accuracy of 96.8%, precision of 97.1%, recall of 96.5%, F1-score of 96.8%, and an ROC-AUC of 0.985, with narrow confidence intervals (CIs) indicating statistical robustness. Comparative analysis with baseline convolutional neural network (CNN) architectures (Standard VGG16, ResNet50, DenseNet121, and VGG16-CAE without BiLSTM) and recent literature demonstrated that the proposed method significantly outperforms existing approaches by effectively capturing sequential dependencies in MRI data. The modular design of the framework allows for future extensions, including tumor localization and segmentation tasks, making it a promising solution for clinical applications. This work highlights the advantage of combining hierarchical feature learning with sequential modeling for precise BT detection, contributing to improved decision support in neuro-oncology.

Keywords

Bi-Long short-term memory, convolutional autoencoder, VGG16 layer, deep learning

1. Introduction

Brain tumors (BTs) are among the most critical and life-threatening medical conditions, posing significant challenges in early diagnosis, treatment planning, and monitoring. Accurate and timely detection

of BTs is essential, as it directly affects patient prognosis, treatment decisions, and overall survival rate. Magnetic resonance imaging (MRI) has long been the imaging modality of choice for BT diagnosis due to its superior contrast resolution, ability to visualize soft tissues, and non-invasive nature [1–4]. However,

manual inspection of MRI scans by radiologists is time-consuming, subjective, and prone to human error, especially when dealing with large volumes of data or subtle tumor manifestations. Therefore, the need for automated, reliable, and efficient computational frameworks for BT detection has become increasingly important in modern medical diagnostics [5–8]. In recent years, the rapid advancement of deep learning techniques has provided promising solutions for automated medical image analysis, leveraging powerful models that can automatically extract relevant features from high-dimensional imaging data and perform accurate classification or detection tasks. Convolutional neural networks (CNNs), in particular, have been widely adopted for medical image classification, detection, and segmentation tasks, as they can effectively capture spatial hierarchies in images [9–12]. Among various CNN architectures, VGG16 has emerged as one of the most popular models due to its simple yet effective design based on stacked convolutional layers and its ability to extract deep spatial features. VGG16 has been widely used in medical image classification tasks, including BT classification, due to its strong feature extraction capability [13–17]. However, CNN models typically focus on spatial features and may not explicitly capture sequential dependencies or feature correlations across layers, which can limit their performance in capturing more complex representations present in medical imaging datasets [18–22]. To overcome these limitations, sequential modeling techniques, such as recurrent neural networks (RNNs), and particularly Bi-Directional Long Short-Term Memory (BiLSTM) networks, can significantly enhance feature representation by learning temporal dependencies in both forward and backward directions. BiLSTMs have demonstrated remarkable success in sequence modeling tasks such as speech recognition, natural language processing, and time-series forecasting, but their application in medical image classification remains relatively underexplored. This research proposes a novel deep learning framework that combines a VGG16-based Convolutional Autoencoder (CAE) with a BiLSTM network for BT classification from MRI images. The proposed framework is designed to leverage the powerful feature extraction capability of the VGG16-based autoencoder, which reduces the dimensionality of the input while preserving essential features, and the BiLSTM network, which refines the extracted features by capturing sequential dependencies to improve classification accuracy. In contrast to traditional CNN-based classifiers, which treat each sample independently, the proposed method allows

for learning contextual information across the feature space, enhancing its discriminative power.

II. Literature Review

BT detection and classification using MRI images have become significant research areas in medical imaging due to their potential to improve diagnostic accuracy and assist clinicians in treatment planning. In recent years, deep learning methods, particularly CNNs, have demonstrated remarkable success in automating feature extraction and classification tasks from high-dimensional medical images, outperforming traditional machine learning techniques. Shib et al. [1] proposed a deep learning framework utilizing VGG16 for BT detection and classification, demonstrating that pretrained VGG16 models, when fine-tuned on brain MRI datasets, could significantly improve classification accuracy compared with baseline CNN models. However, their approach primarily focused on supervised classification without integrating temporal dependencies or sequential feature refinement, which limits its ability to capture complex spatial feature interactions across layers.

Hafeez et al. [2] developed a CNN-based model for BT classification using MRI images and reported promising accuracy levels. Their study reinforced the importance of convolutional feature extraction for high-dimensional medical images but similarly lacked exploration of advanced sequential modeling techniques such as recurrent architectures, which could improve learning contextual dependencies across features. In another study, Manjunath et al. [6] investigated fine-tuning deep learning models for BT classification, showing that pretrained architectures provide a strong baseline for performance, but still required extensive hyperparameter tuning and did not evaluate hybrid models combining different deep learning paradigms.

Khan et al. [4] introduced a hybrid network model, Hybrid-Net, combining DenseNet-169 with traditional machine learning classifiers, such as support vector machines (SVMs), to enhance classification performance in BT diagnosis. Their study emphasized the complementary strengths of deep convolutional feature extraction and classical classifiers, but lacked a systematic ablation study to quantify the individual contributions of each component, nor did it explore temporal modeling layers like BiLSTM.

From a data optimization perspective, Gowda and Lakshmikantha [5] proposed hybrid swarm optimization algorithms for clustering streaming data, highlighting the importance of adaptive optimization methods

to handle large-scale and continuous data in medical applications. Although not specific to BT classification, such adaptive optimization approaches inspire future directions in model optimization for evolving datasets. Similarly, Manjunath et al. [6] focused on automated liver segmentation using a modified ResUNet model on CT images, underscoring the growing trend of encoder-decoder architectures for accurate feature extraction and segmentation in medical imaging. Although their work targeted a different organ, the principle of automated feature extraction remains highly relevant to tumor classification tasks.

Bogacsovics et al. [7] proposed diverse ensemble architectures for automatic BT classification, where multiple deep learning models were combined to improve robustness and generalization. Their ensemble approach achieved higher accuracy by fusing the outputs of various classifiers, but the study did not investigate sequential modeling or temporal dependencies that could further improve feature extraction and classification. In a complementary work, Saeedi et al. [8] combined convolutional deep learning methods with machine learning classifiers for BT detection, reporting significant improvements in classification accuracy. However, their approach followed a two-step pipeline of feature extraction followed by classification, which does not allow joint optimization of both tasks in an end-to-end trainable framework.

Bouhafra and El Bahi [9] conducted a comprehensive systematic review of deep learning methods for BT detection and classification using MRI images published between 2020 and 2024. They concluded that convolutional architectures such as VGG16, ResNet, and DenseNet dominate the field due to their strong spatial feature extraction abilities. Nevertheless, they highlighted a persistent gap in the literature regarding temporal feature modeling and unified frameworks that integrate feature extraction and sequential modeling within a single architecture. In parallel, Alrashedy et al. [10] proposed BrainGAN, a framework using generative adversarial networks (GANs) combined with CNN classifiers for synthetic MRI image generation and tumor classification. Their work illustrated that synthetic data augmentation can help mitigate the limited size of annotated medical datasets, but did not apply recurrent architectures to capture temporal dependencies between image features.

More recent advances include transformer-based methods, which provide powerful attention mechanisms for feature modeling. Tabatabaei et al. [11]

introduced a fusion-based deep learning architecture combining attention transformers with CNNs for MRI BT classification, significantly improving feature interpretability and classification accuracy. Their approach highlights the benefit of attention mechanisms in focusing on relevant spatial regions of MRI images but remains computationally intensive and lacks explicit temporal feature modeling. Similarly, Şahin et al. [12] investigated multiobjective optimization of the vision transformer (ViT) architecture for efficient BT classification, focusing on optimizing model size, accuracy, and inference speed. While the study provided strong evidence for transformer-based models, the reliance on global self-attention mechanisms may overlook local sequential dependencies that BiLSTM layers can naturally model.

While previous research has demonstrated the effectiveness of CNNs such as VGG16, DenseNet, and ResNet for BT classification, there is a notable research gap in integrating convolutional feature extraction with sequential learning mechanisms such as BiLSTM to capture both spatial and temporal dependencies within the data. Most existing works focus on either spatial feature extraction or ensemble learning without systematically studying the contribution of each component. Furthermore, many studies lack rigorous hyperparameter tuning strategies and cross-validation techniques to improve reproducibility and generalization. This research seeks to fill these gaps by proposing a deep learning framework that integrates a VGG16-based CAE with BiLSTM for enhanced BT classification, supported by a detailed ablation study, hyperparameter optimization, and robust cross-validation approach. Table 1 shows a summary of the existing studies discussed in the literature section.

III. Methodology

a. Dataset description and source

The proposed framework for BT classification is evaluated using the publicly available Kaggle Brain MRI Images Dataset [23]. The dataset consists of MRI images of brain scans labeled as either “Tumor” or “No Tumor.” Specifically, the dataset contains approximately 3,000 images divided into two categories:

- Tumor Images: 1,500 images
- No Tumor Images: 1,500 images

Table 1: Comparative analysis of recent BT classification studies

Study	Methodology	Dataset	Key contributions	Limitations
Shib et al. [1]	VGG16-based CNN	Kaggle Brain MRI	Demonstrated pretrained CNNs improve classification accuracy	No temporal/sequential modeling; single-step classification
Hafeez et al. [2]	CNN with custom layers	Public MRI dataset	Reinforced importance of convolutional feature extraction	No recurrent or sequential modeling explored
Manjunath et al. [6]	Fine-tuned deep learning models	Brain MRI Dataset	Pretrained architectures provide strong baseline	Extensive hyperparameter tuning required; no hybrid model
Khan et al. [4]	Hybrid-Net (DenseNet + SVM)	Public MRI dataset	Combined deep features with traditional classifiers	No ablation study; lacks temporal modeling
Saeedi et al. [8]	CNN + SVM	Public Brain MRI	Two-step pipeline (feature extraction + classification)	No end-to-end joint optimization
Bouhafra & El Bahi [9]	Systematic review	Multiple MRI datasets	Comprehensive survey of deep learning methods	No sequential modeling in proposed frameworks
Alrashedy et al. [10]	BrainGAN (GAN + CNN)	Public brain MRI	Synthetic data augmentation to address dataset size	No recurrent/temporal modeling
Tabatabaei et al. [11]	CNN + transformer	Brain MRI Dataset	Attention-based feature modeling improves interpretability	High computational cost; lacks explicit temporal modeling
Şahin et al. [12]	ViT	Brain MRI Dataset	Multiobjective optimization for efficient classification	Global attention may miss local sequential dependencies

BT, brain tumor; CNNs, convolutional neural networks; GAN, generative adversarial network; MRI, magnetic resonance imaging; SVM, support vector machine; ViT, vision transformer.

Each image is provided in grayscale format with a resolution of 256×256 pixels. The dataset includes images collected from multiple medical centers, making it representative of real-world variability in brain MRI scans. Prior to model training, standard preprocessing steps such as resizing, normalization (scaling pixel values to $[0, 1]$), and data augmentation (rotation, flipping, scaling) are applied to increase dataset diversity and mitigate overfitting.

To ensure robust evaluation and generalization of the model, a fivefold cross-validation strategy is adopted. This divides the dataset into five subsets, where in each fold, four subsets are used for training and one for validation, rotating in each iteration.

b. Architecture of predicting BT diagnosis using BiLSTM with VGG16

The proposed deep learning architecture for BT classification integrates a VGG16-based CAE with

a BiLSTM network, forming an end-to-end trainable system that extracts spatial features from MRI images and captures sequential dependencies in the feature space. The architecture consists of two key components: the CAE serves as a deep feature extractor, while the BiLSTM module refines and models the extracted features for final classification. The design leverages the hierarchical feature learning capability of VGG16 and the sequential learning capability of BiLSTM to enhance classification accuracy and robustness.

b.i Input preprocessing

Let the input dataset $D = \{(X_i, Y_i)\}$ consist of N MRI images $X_i \in \mathbb{R}^{H \times W \times C}$ and their corresponding labels $Y_i \in \{0, 1\}$, where 0 denotes a normal brain and 1 denotes the presence of a tumor. Each input image X_i is first resized to $224 \times 224 \times 3$ to match the input dimension of VGG16. Grayscale images are

replicated across three channels to form a 3-channel image. The preprocessing step also includes normalization of pixel values to $[0, 1]$, denoted as follows:

$$X_i = \frac{X_i}{255}$$

Data augmentation, such as rotation, flipping, and scaling, is applied to increase the diversity of the dataset and reduce overfitting.

b.ii VGG16-based CAE

The first stage of the architecture employs a CAE with VGG16 as the encoder. Autoencoders are designed to learn a compressed representation of the input while preserving essential features. In this framework, the encoder leverages the pretrained VGG16 network (excluding the fully connected top layers) to extract high-dimensional spatial features from MRI images. Figure 1 shows the architecture of the proposed work.

Let the encoder function be represented as follows:

$$F_i = f_{\{VGG16-Encoder\}}(X_i, \mathcal{O}_e)$$

where $F_i \in R^{h \times w \times d}$ is the feature map generated by the encoder, and $\mathcal{O}_e$ denotes the learnable parameters of the encoder. For an input size of $224 \times 224 \times 3$, the output feature map typically has reduced spatial dimensions $h = w = 7$ and depth $d = 512$. This feature map preserves hierarchical spatial information while reducing the dimensionality of the input data.

The decoder reconstructs the input from the feature map using upsampling and convolutional layers:

$$X_i = f_{\text{decoder}}(F_i, \mathcal{O}_d)$$

where $X_i \in R^{224 \times 224 \times 3}$ is the reconstructed image and $\mathcal{O}_d$ represents the decoder parameters. Although reconstruction is part of the training process, the primary output used for classification is the latent feature representation F_i .

To convert the feature map F_i into a suitable format for sequential modeling with BiLSTM, it is flattened as follows:

$$V_i = \text{Flatten}(F_i) \in R^{h \times w \times d}$$

where V_i is the one-dimensional feature vector for the i -th image.




Figure 1: Overview of the architecture. MRI, magnetic resonance imaging.

b.iii BiLSTM

After feature extraction, the flattened feature vector V_i is fed into a BiLSTM network. LSTM networks are designed to capture long-range dependencies in sequential data by maintaining a memory cell C_t and hidden state h_t at each time step t . The LSTM cell can be mathematically formulated as follows:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f)$$

$$i_t = \sigma(W_i \cdot [h_t - 1, x_t] + b_i)$$

$$C_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c)$$

$$O_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o)$$

$$h_t = o_t \odot \tanh(C_t)$$

where σ denotes the sigmoid activation, $\odot$ denotes element-wise multiplication, and W , b are learnable parameters.

In BiLSTM, two LSTM layers process the input sequence in forward and backward directions:

$$h_t = LSTM_{\text{forward}}(V_i)$$

$$h_t = LSTM_{\text{backward}}(V_i)$$

The outputs from both directions are concatenated as follows:

$$h_t^{b_i} = [h_t; h_t]$$

This concatenated representation captures dependencies from both past and future contexts within the feature sequence, enhancing discriminative capability for classification.

Feature reshaping and sequential modeling justification

The latent feature map output by the VGG16 encoder has a spatial dimension of 7×7 with 512 channels, resulting in a tensor of shape (7, 7, 512). To feed this tensor into the BiLSTM, it is reshaped into a sequence of 49 steps (7×7), where each step corresponds to a 512-dimensional feature vector as follows:

$$V_i = \text{Reshape}(F_i) \in R^{(49, 512)}$$

Here, each “time step” of the sequence represents a local spatial patch of the original image, preserving hierarchical feature information across spatial locations. By interpreting the flattened spatial patches as sequential inputs, the BiLSTM can model long-range dependencies between spatially distant regions in

the MRI, which is particularly useful for capturing structural correlations and contextual relationships of tumor regions within the brain. This sequential modeling allows the network to leverage both forward and backward dependencies, enhancing the discriminative power of the classifier beyond conventional feed-forward CNNs.

Rationale for using BiLSTM in static MRI image classification

Although MRI images are static in nature, the feature representations extracted from the VGG16-based encoder contain structured spatial dependencies across pixels and regions that exhibit sequential-like correlations. When the 2D feature map ($7 \times 7 \times 512$) is flattened into a feature vector, the spatial arrangement of these features can be interpreted as a structured sequence of local descriptors. The BiLSTM network is employed not to capture temporal dynamics in time, but to model contextual relationships between these ordered spatial features.

This approach allows the model to learn how distant spatial regions within the MRI contribute jointly to tumor characterization—something that standard CNNs or fully connected layers cannot effectively represent due to their limited receptive context. The bidirectional mechanism of BiLSTM captures dependencies in both forward and reverse directions across the spatial sequence, providing a more holistic contextual understanding of anatomical patterns and texture gradients.

Empirically, the inclusion of BiLSTM improves performance by approximately 2.5% in accuracy compared with the VGG16-CAE-only model (Table 5). This demonstrates that modeling inter-feature dependencies enhances discriminative power, even in static image domains, by reinforcing relationships among anatomically correlated structures. Such hybrid spatial-contextual modeling has been successfully validated in other medical imaging works as well, where recurrent networks improve interpretability and feature coherence across complex spatial patterns.

b.iv Classification layer

The BiLSTM output $h_t^{b_i}$ is passed through a fully connected dense layer followed by a Softmax activation to generate class probabilities:

$$Y_i = \text{Softmax}(W_c \cdot h_t^{b_i} + b_c)$$

where $Y_i \in [0,1]^2$ represents the predicted probability for each class (Tumor/No Tumor), W_c and b_c are learnable weights of the dense layer. The predicted label is obtained as follows:

$$Y_i^{\text{pred}} = \text{argmax}(Y_i)$$

The model is trained end-to-end by minimizing the ategorical cross-entropy loss:

$$L = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^2 Y_{i,c} \log(Y_{i,c})$$

where $Y_{i,c}$ is the ground truth label for class c .

c. Working of CAE using VGG16

The CAE is designed to compress the input MRI images into a lower-dimensional representation (latent features) while preserving important spatial information. The encoder component of the CAE uses a pretrained VGG16 network (without top dense layers) as the backbone to extract hierarchical spatial features.

c.i Encoder (VGG16-based)

- The input image of size $256 \times 256 \times 1$ is first resized to match VGG16 input dimensions ($224 \times 224 \times 3$) by duplicating the grayscale channel.
- The convolutional layers of VGG16 process the input image, passing through multiple convolution and max-pooling layers to extract rich hierarchical features.
- The output of the encoder is a feature map of reduced spatial dimensions but high feature depth (e.g., $7 \times 7 \times 512$).

c.ii Decoder (reconstruction layer)

- The decoder reconstructs the image from the latent feature representation using upsampling and convolution layers.
- While the primary purpose of the decoder is reconstruction, in this framework, the latent feature vector (output of the encoder) is used for classification downstream.

By using VGG16 as the encoder, the model benefits from pretrained weights (on ImageNet) that help in better generalization and fast convergence, especially with limited medical data.

d. Working of BiLSTM

After the feature extraction via the VGG16-based encoder, the latent feature map is flattened into a feature vector (e.g., $7 \times 7 \times 512 \rightarrow 25,088$ -dimensional vector).

The BiLSTM module is used to model sequential dependencies across the feature vector. The core idea is to treat the flattened feature as a temporal sequence where each feature corresponds to a time step. The BiLSTM consists of the following:

- Two LSTM layers: one processes the sequence from start to end (forward direction), and the other processes it in reverse (backward direction).
- The outputs from both directions are concatenated, providing a comprehensive understanding of sequential feature dependencies.

The BiLSTM captures contextual relationships between features that may represent structural dependencies in the MRI image, enhancing discriminative power.

The final hidden states are passed through a fully connected (Dense) layer followed by a Softmax activation function to produce probabilistic classification outputs (Tumor/No Tumor).

Algorithm: Brain Tumor Diagnosis using VGG16-based CAE and BiLSTM

Input:

- MRI Dataset $D = \{X_i, Y_i\}$ where $X_i \in \mathbb{R}^{(256 \times 256 \times 1)}$, $Y_i \in \{0,1\}$

Output:

- Predicted Class (Tumor or No Tumor)

Steps:

1. Preprocessing:
 - Resize X_i to $224 \times 224 \times 3$
 - Normalize pixel values to $[0,1]$
 - Apply data augmentation (rotation, flipping)
2. Feature Extraction using VGG16-CAE:
 - Pass input X_i through VGG16 Encoder (without top layers)
 - Obtain feature map $F_i \in \mathbb{R}^{(7 \times 7 \times 512)}$
3. Flatten Features:
 - Flatten F_i to create a feature vector $V_i \in \mathbb{R}^{25088}$

4. Sequential Feature Modeling using BiLSTM:
 - Input V_i into BiLSTM layers:
 - Forward LSTM: Processes V_i from first to last element
 - Backward LSTM: Processes V_i in reverse
 - Concatenate outputs of both directions
5. Classification:
 - Pass BiLSTM output to Dense layer
 - Apply Softmax activation to obtain classification probabilities $P(Y_i)$
6. Training:
 - Loss Function: Categorical Cross-Entropy
 - Optimizer: Adam with tuned learning rate
 - Hyperparameters: Batch size, number of BiLSTM units, epochs (optimized by grid search)
7. Validation:
 - Apply 5-fold Cross-Validation
 - Compute metrics: Accuracy, Precision, Recall, F1-score
8. Output:
 - Final Prediction: $Y_{pred} \in \{0:NoTumor, 1:Tumor\}$

End Algorithm

The proposed framework leverages a VGG16-based CAE to ensure effective hierarchical spatial feature extraction from MRI images, enabling the model to capture both low-level and high-level features that are critical for accurate BT classification. The use of a pretrained VGG16 backbone allows the CAE to extract robust representations of anatomical structures, capturing edges, textures, and complex spatial patterns that are often challenging to discern in raw MRI images. By compressing the high-dimensional input into a meaningful latent feature space, the CAE not only reduces computational complexity but also preserves essential information necessary for distinguishing tumor regions from normal tissue. To complement this spatial feature extraction, the framework integrates a BiLSTM module, which further enhances feature representation by modeling long-range dependencies between sequential features in the latent space. The BiLSTM processes the flattened feature vectors in both forward and backward directions, allowing the network to capture contextual relationships across feature dimensions that are otherwise missed by conventional feedforward architectures. This sequential modeling capability improves the discriminative power of the network,

making it more sensitive to subtle variations in tumor morphology and structure. In addition to architectural design, systematic hyperparameter tuning is employed to guarantee reproducibility and stability of the model, including optimization of parameters such as learning rate, batch size, number of BiLSTM units, and dropout rates. Careful selection of these parameters ensures that the network converges efficiently while minimizing overfitting, and provides consistent performance across multiple training runs. To further enhance the model's generalizability, a fivefold cross-validation strategy is implemented, in which the dataset is divided into five subsets, and the model is trained and validated on different combinations of these subsets. This procedure not only provides a reliable estimate of the model's performance on unseen data but also mitigates the risk of overfitting that can occur when using a single train-test split, especially given the limited size of publicly available brain MRI datasets. The modular design of the framework allows for flexibility and scalability, enabling easy integration of additional components such as tumor localization, segmentation modules, or attention mechanisms in future extensions. By decoupling the feature extraction and sequential modeling stages, the architecture can be adapted to different medical imaging tasks or combined with other deep learning techniques, making it a versatile tool for both research and clinical applications. Overall, the proposed approach effectively combines hierarchical spatial feature learning, sequential dependency modeling, rigorous hyperparameter optimization, and cross-validation strategies, resulting in a robust, generalizable, and extensible framework for BT classification from MRI images, while also providing a foundation for future developments in automated tumor localization and segmentation.

d.i Cross-validation for robust evaluation

To address concerns regarding model generalizability and avoid bias introduced by a fixed 88%–12% train-test split, the proposed framework employs five-fold cross-validation. In this approach, the dataset is partitioned into five equally sized folds. During each iteration, one fold is used as the validation set while the remaining four folds are used for training. This process is repeated five times, ensuring that every sample in the dataset is used once for validation and four times for training. The final performance metrics—Accuracy, Precision, Recall, F1-Score, and ROC-AUC—are averaged across all folds to provide a more reliable estimate of model performance.

Formally, let the dataset be $D = \{(x_i, y_i)\}_{i=1}^N$ with N samples. For fold $k \in \{1, 2, 3, 4, 5\}$

Training set $D_{\text{train}}^k = D \setminus D_{\text{val}}^k$, Validation set D_{val}^k

where D_{val}^k is the k -th fold of size $N/5$. The model is trained on D_{train}^k and evaluated on D_{val}^k . Let M denote a performance metric (e.g., Accuracy). The cross-validated performance is computed as follows:

$$M = \frac{1}{K} \sum_{k=1}^K M(D_{\text{val}}^k) \text{ Standard Deviation:}$$

$$\sigma M = \sqrt{\frac{1}{K} \sum_{k=1}^K (M(D_{\text{val}}^k) - M)^2}$$

where $K = 5$.

This method ensures that performance evaluation is not biased by a single train-test split, providing a more accurate estimate of the model's generalization capability. Compared with the previously used 88%–12% split, fivefold cross-validation reduces overfitting risk and allows the calculation of confidence intervals (CIs) and standard deviations for all metrics, providing statistical rigor and reproducibility. The results from the fivefold cross-validation confirm the *robustness and stability* of the proposed VGG16-CAE + BiLSTM framework, as evidenced by narrow CIs and low standard deviations across Accuracy, Precision, Recall, F1-Score, and ROC-AUC.

IV. Experimental Setup and Results

The proposed VGG16-based CAE with BiLSTM framework was implemented using Python 3.10, TensorFlow 2.12, and Keras. Experiments were conducted on a workstation equipped with an Intel i9 CPU, 32 GB RAM, and NVIDIA RTX 3090 GPU to ensure efficient processing of large MRI datasets.

a. Data partitioning and validation strategy

The Kaggle Brain MRI dataset used in this study consists of approximately 3,000 T1-weighted MRI images categorized into tumor and non-tumor classes. To mitigate the risk of data leakage, a patient-wise data separation protocol was implemented during dataset partitioning. Specifically, all MRI slices corresponding to a single patient were allocated exclusively to one subset—training, validation, or testing—ensuring that no patient's data appeared across multiple folds. This strategy maintains strict independence between

samples used for model optimization and those used for evaluation.

Although the dataset is relatively modest in size for deep learning, this limitation was addressed through several measures as follows:

1. **Data augmentation** (including rotation, flipping, contrast normalization, and zooming) was employed to synthetically expand the diversity of the training data;
2. **Transfer learning** from a pretrained VGG16 backbone helped reduce the data requirement by leveraging generalized visual representations;
3. **Fivefold cross-validation** was used to obtain a more stable and unbiased estimate of model performance.

For robustness assessment, the proposed CAE–BiLSTM model was further tested on a held-out subset representing 20% of the total data, never seen during training. In future work, the framework will be validated on larger benchmark datasets such as BraTS 2021 and TCGA-LGG/GBM, to confirm the model's generalizability and clinical utility.

To robustly evaluate model performance and generalizability, we employed fivefold cross-validation, where the dataset was divided into five equal parts. Each fold iteratively serves as the validation set while the remaining four are used for training. This systematic evaluation ensures that the results are not biased by a particular data split and reduces overfitting risk.

b. Hyperparameter tuning

Hyperparameters were tuned using a grid search strategy, systematically varying key parameters to identify optimal values. The final configuration is shown in Table 2.

Early stopping was enabled based on validation loss, stopping training after 10 consecutive epochs without improvement, to prevent overfitting and ensure reproducibility.

The proposed framework was rigorously evaluated using multiple quantitative performance metrics to assess its effectiveness in BT classification. The key evaluation measures include Accuracy, Precision, Recall, F1-Score, and ROC-AUC, which collectively provide a comprehensive assessment of classification quality, reliability, and discriminative power. Importantly, to ensure statistical rigor and reproducibility, the results were averaged over a fivefold cross-validation process, and 95% CIs

were computed for each metric. This approach guarantees that the performance is not biased by a particular data split and provides a measure of variability across different training and validation sets. As shown in Table 3, the proposed model consistently achieves high values across all metrics with narrow CIs, indicating both robustness and stability in tumor classification performance.

These metrics indicate not only high performance but also strong statistical confidence in the results, demonstrating consistency across multiple data splits.

c. Error analysis

The performance stability was assessed through standard deviation calculations over the fivefold cross-validation, as shown in Table 4.

d. Model interpretability

To provide interpretability for the proposed VGG16-CAE + BiLSTM framework, Grad-CAM attention maps were generated to highlight regions of the MRI images that contribute most to the model's predictions.

Table 2: Hyperparameter tuning

Hyperparameter	Final value
Learning rate	0.0001
Optimizer	Adam
Batch size	32
Epochs	50
BiLSTM units	128
Dropout rate	0.5

BiLSTM, Bi-Directional Long Short-Term Memory.

Table 3: Quantitative performance metrics

Metric	Mean (%)	95% CI (%)
Accuracy	96.8	±0.5
Precision	97.1	±0.4
Recall	96.5	±0.6
F1-score	96.8	±0.5
ROC-AUC	0.985	±0.004

CI, confidence interval.

Figure 2 shows representative attention maps over synthetic tumor-like MRI images, where bright regions indicate areas of high activation, highlighting tumor regions in the MRI images. These visualizations demonstrate that the model focuses on clinically relevant areas for BT classification, providing interpretability and supporting the reliability of predictions. The aggregated confusion matrix across all folds provides a clear view of the classification performance as in Table 5.

Figure 3 highlights the model's effectiveness in minimizing false positives and false negatives, crucial for clinical applications.

e. Training stability and convergence analysis

To evaluate the training stability and monitor overfitting, the model's training and validation loss and accuracy curves were plotted across epochs. Figure 4 illustrates that the proposed VGG16-CAE + BiLSTM framework converges steadily, with both training and validation loss decreasing consistently while accuracy increases. Early stopping with a patience of 10 epochs and dropout regularization (rate = 0.5) were applied during training to prevent overfitting. The minimal gap between training and validation curves confirms that the model generalizes well and that overfitting is effectively controlled. These convergence plots provide strong evidence of the robustness and stability of the proposed framework.

The convergence of both loss and accuracy demonstrates stable training, and the minimal gap between training and validation curves indicates effective prevention of overfitting through early stopping and dropout regularization.

To further validate the effectiveness of the proposed VGG16-CAE + BiLSTM framework, a comparative evaluation as in Table 6 was conducted against several well-established baseline deep learning models commonly used in medical image classification tasks. These baseline models include Standard VGG16, ResNet50, DenseNet121, and VGG16-CAE without BiLSTM. Each model was trained and

Table 4: Error analysis

Metric	Standard deviation (%)
Accuracy	0.22
Precision	0.18
Recall	0.24
F1-score	0.21




Figure 2: Grad-CAM attention maps generated from the VGG16 encoder of the proposed VGG16-CAE + BiLSTM framework. BiLSTM, Bi-Directional Long Short-Term Memory; CAE, convolutional autoencoder.

Table 5: Confusion matrix

	Predicted tumor	Predicted no tumor
Actual tumor	291 $\pm$ 5	9 $\pm$ 4
Actual no tumor	11 $\pm$ 3	289 $\pm$ 5

evaluated using the same dataset and experimental setup, including identical preprocessing,

hyperparameter tuning strategies, and fivefold cross-validation, to ensure a fair and consistent comparison.

Figure 5 clearly demonstrates that integrating BiLSTM improves performance significantly over the CAE-only model and standard CNNs.

The ROC curve analysis confirms superior class separation capabilities of the proposed framework as in Table 7, with the highest AUC compared with baseline models.




Figure 3: Confusion matrix.




Figure 4: Training and validation loss and accuracy curves for the VGG16-CAE + BiLSTM framework. BiLSTM, Bi-Directional Long Short-Term Memory; CAE, convolutional autoencoder.

Table 6: Comparative evaluation of the proposed framework against baseline models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
Standard VGG16 (pretrained CNN)	91.2 $\pm$ 0.8	91.5 $\pm$ 0.7	90.8 $\pm$ 0.9	91.1 $\pm$ 0.8	0.943 $\pm$ 0.007
ResNet50 (pretrained CNN)	92.5 $\pm$ 0.7	92.7 $\pm$ 0.6	92.0 $\pm$ 0.8	92.3 $\pm$ 0.7	0.951 $\pm$ 0.006
DenseNet121 (pretrained CNN)	93.1 $\pm$ 0.7	93.5 $\pm$ 0.6	92.7 $\pm$ 0.7	93.1 $\pm$ 0.7	0.957 $\pm$ 0.005
VGG16-CAE (without BiLSTM)	94.3 $\pm$ 0.6	94.5 $\pm$ 0.5	94.0 $\pm$ 0.6	94.2 $\pm$ 0.5	0.970 $\pm$ 0.005
Proposed VGG16-CAE + BiLSTM	96.8 $\pm$ 0.5	97.1 $\pm$ 0.4	96.5 $\pm$ 0.6	96.8 $\pm$ 0.5	0.985 $\pm$ 0.004

BiLSTM, Bi-Directional Long Short-Term Memory; CAE, convolutional autoencoder; CNNs, convolutional neural networks; ROC-AUC.




Figure 5: Comparison of model performance across various models.

To place the performance of the proposed VGG16-CAE + BiLSTM framework in the context of existing research, a comparative analysis was conducted against several recent studies from the literature. These studies represent a range of methodologies employed for BT detection and classification, including pure CNN-based models, hybrid CNN + traditional classifier systems, and fine-tuned deep learning models. The comparison considers metrics such as Accuracy, Precision, Recall, and F1-Score, as well as key observations regarding the architectural design of each study. The objective is to highlight the distinct

advantages of the proposed framework, particularly its integration of hierarchical spatial feature extraction (via VGG16-CAE) and long-range sequential dependency modeling (via BiLSTM), which contribute to enhanced tumor classification performance. As shown in Table 8, the proposed method achieves the highest accuracy and F1-score among the compared studies, supported by systematic hyperparameter tuning and cross-validation. This confirms the model's superior ability to generalize and capture complex patterns in MRI data, offering significant improvements over prior approaches that do not model feature dependencies

explicitly. Figures 6–9 show the performance of the proposed model with various metrics.

f. Ablation study

To validate the contribution of each component, an ablation study was conducted by removing or replacing key modules:

- 1. VGG16-CAE only (without BiLSTM): Accuracy drops to 94.3%, confirming the importance of sequential feature modeling.
- 2. BiLSTM with raw pixels (without CAE): Accuracy reduces to 92.0%, showing the necessity of hierarchical feature extraction.
- 3. Standard VGG16 CNN: Accuracy is 91.2%, demonstrating the advantage of CAE for deep feature representation.
- 4. Proposed VGG16-CAE + BiLSTM: Achieves 96.8%, highlighting the synergistic effect of combining CAE and BiLSTM.

Table 7: ROC analysis

Model	ROC-AUC
VGG16	0.943
ResNet50	0.951
DenseNet121	0.957
VGG16-CAE (without BiLSTM)	0.970
Proposed VGG16-CAE + BiLSTM	0.985

BiLSTM, Bi-Directional Long Short-Term Memory; CAE, convolutional autoencoder.

The ablation results confirm that both CAE and BiLSTM modules are crucial for achieving optimal performance, with BiLSTM providing sequential context and CAE extracting hierarchical spatial features.

The results confirm that the proposed framework significantly outperforms baseline CNN models and prior deep learning-based BT detection methods in the literature. The integration of hierarchical spatial feature extraction (via VGG16-CAE) and sequential dependency modeling (via BiLSTM) leads to superior classification performance with high accuracy (96.8%), robust precision, recall, and F1-score, and excellent class separation (ROC-AUC of 0.985). Systematic hyperparameter tuning, fivefold cross-validation, and CI analysis provide strong statistical evidence of reproducibility and generalization.

V. Conclusion

In this study, we proposed a robust and effective deep learning framework for automated BT detection from MRI images, integrating a VGG16-based CAE with a BiLSTM network. The VGG16-based CAE enabled efficient hierarchical spatial feature extraction from MRI images, while the BiLSTM module successfully captured long-range dependencies and contextual relationships between features, enhancing the model’s ability to discriminate between tumor and non-tumor cases. Extensive experiments were conducted using the publicly available Kaggle Brain MRI Dataset, with systematic hyperparameter tuning and fivefold cross-validation

Table 8: Comparative analysis of the proposed framework with recent literature

Study	Methodology	Dataset	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Shib et al. [1]	VGG16-based CNN	Brain MRI Kaggle Dataset	94.0	94.3	93.8	94.1
Hafeez et al. [2]	CNN with custom layers	Public MRI dataset	92.5	93.0	91.8	92.4
Manjunath et al. (2025) [6]	Fine-tuned deep learning models	Brain MRI Dataset	95.1	95.3	94.8	95.0
Saeedi et al. (2023) [8]	CNN + SVM Classifier	Public brain MRI Dataset	93.7	94.0	93.2	93.6
Proposed VGG16-CAE + BiLSTM	VGG16-based CAE + BiLSTM	Kaggle Brain MRI Dataset	96.8 ± 0.5	97.1 ± 0.4	96.5 ± 0.6	96.8 ± 0.5

BiLSTM, Bi-Directional Long Short-Term Memory; CAE, convolutional autoencoder; CNNs, convolutional neural networks; MRI, magnetic resonance imaging; SVMs, support vector machines.




Figure 6: Accuracy analysis.




Figure 7: Precision analysis.

ensuring reproducibility and generalization. The proposed model achieved outstanding performance, with an average accuracy of 96.8%, precision of 97.1%, recall of 96.5%, and F1-score of 96.8%, along with a high ROC-AUC of 0.985. Statistical rigor was incorporated by reporting CIs and standard deviations, confirming consistent performance across data splits. A comprehensive comparative

analysis with both baseline models (Standard VGG16, ResNet50, DenseNet121, and VGG16-CAE without BiLSTM) and the recent literature demonstrated that the proposed framework significantly outperforms existing methods, particularly due to the explicit sequential feature modeling enabled by BiLSTM. This approach addresses the limitations of prior works that relied solely on CNN-based feature




Figure 8: Recall analysis.




Figure 9: F1 score comparison.

extraction without modeling sequential dependencies, proving its superior ability to capture complex patterns in MRI data. Furthermore, the proposed framework is modular and scalable, providing a solid foundation for future research. Potential extensions include integrating tumor localization or segmentation modules, further improving diagnostic precision, and making the system suitable for end-to-end clinical applications.

References

[1] Shib, S. K., Dipto, D. R., Rahman, M. T., & Shufian, A. (2025). Deep learning-based brain tumor detection and classification using VGG16. In 2025 4th International Conference on Robotics, Electrical and Signal Processing Techniques (ICREST) (pp. 482–486). Dhaka, Bangladesh. <https://doi.org/10.1109/ICREST63960.2025.10914441>

- [2] Hafeez, M. A., Kayasandik, C. B., & Dogan, M. Y. (2022). Brain tumor classification using MRI images and convolutional neural networks. In 2022 30th Signal Processing and Communications Applications Conference (SIU) (pp. 1–4). <https://doi.org/10.1109/SIU55643.2022.9876543>
- [3] V., R., M., N., Y. G., H. M., M., et al. (2025). Brain tumor classification on MRI images using fine-tuned deep learning models. *Machine Learning in Computer Science and Engineering*, 1, 22. <https://doi.org/10.1007/s44379-025-00024-0>
- [4] Khan, S. U. R., Zhao, M., Asif, S., & Chen, X. (2023). Hybrid-net: A fusion of DenseNet-169 and advanced machine learning classifiers for enhanced brain tumor diagnosis. *International Journal of Imaging Systems and Technology*, 34, e22975.
- [5] Gowda, Y., & Lakshmikantha, B. R. (2023). On fly hybrid swarm optimization algorithms for clustering of streaming data. *Results in Control and Optimization*, 10, 100190. <https://doi.org/10.1016/j.rico.2022.100190>
- [6] Manjunath, R. V., Gowda, Y., & Manu, H. M. (2025). Automated liver segmentation from CT images using modified ResUNet. *Gastroenterology & Endoscopy*, 3(2), 93–104. <https://doi.org/10.1016/j.gande.2025.03.002>
- [7] Bogacsovics, G., Harangi, B., & Hajdu, A. (2024). Developing diverse ensemble architectures for automatic brain tumor classification. *Multimedia Tools and Applications*. <https://doi.org/10.1007/s11042-024-19657-z>
- [8] Saeedi, S., Rezayi, S., Keshavarz, H., & NiakanKalhori, R. (2023). MRI-based brain tumor detection using convolutional deep learning methods and chosen machine learning techniques. *BMC Medical Informatics and Decision Making*, 23(1), 16. <https://doi.org/10.1186/s12911-023-02114-6>
- [9] Bouhafr, S., & El Bahi, H. (2025). Deep learning approaches for brain tumor detection and classification using MRI images (2020 to 2024): A systematic review. *Journal of Digital Imaging*, 38, 1403–1433. <https://doi.org/10.1007/s10278-024-01283-8>
- [10] Alrashedy, H. H. N., et al. (2022). BrainGAN: Brain MRI image generation and classification framework using GAN architectures and CNN models. *Sensors*, 22(11), 4297. <https://doi.org/10.3390/s22114297>
- [11] Tabatabaei, S., Rezaee, K., & Zhu, M. (2023). Attention transformer mechanism and fusion-based deep learning architecture for MRI brain tumor classification system. *Biomedical Signal Processing and Control*, 86, 105119. <https://doi.org/10.1016/j.bspc.2023.105119>
- [12] Şahin, E., Özdemir, D., & Temurtaş, H. (2024). Multi-objective optimization of ViT architecture for efficient brain tumor classification. *Biomedical Signal Processing and Control*, 91, 105938. <https://doi.org/10.1016/j.bspc.2024.105938>
- [13] Aamir M, Rahman Z, Abro WA, Bhatti UA, Dayo ZA, Ishfaq M. Brain tumor classification utilizing deep features derived from high-quality regions in MRI images. *Biomed Signal Process Control*. 2023;85: 104988. <https://doi.org/10.1016/j.bspc.2023.104988>.
- [14] Kienzler JC, Becher B. Immunity in malignant brain tumors: tumor entities, role of immunotherapy, and specific contribution of myeloid cells to the brain tumor microenvironment. *Eur J Immunol*. 2024;54(2):2250257. <https://doi.org/10.1002/eji.202250257>.
- [15] Feng Y, Cao Y, An D, Liu P, Liao X, Yu B. DAUnet: A U-shaped network combining deep supervision and attention for brain tumor segmentation. *Knowl-Based Syst*. 2024;285: 111348. <https://doi.org/10.1016/j.knosys.2023.111348>.
- [16] Cekic E, Pinar E, Pinar M, Dagcinar A. Deep learning-assisted segmentation and classification of brain tumor types on magnetic resonance and surgical microscope images. *World Neurosurg*. 2024;182:e196–204. <https://doi.org/10.1016/j.wneu.2023.11.073>.
- [17] Zulfiqar F, Bajwa UI, Mehmood Y. Multi-class classification of brain tumor types from MR images using EfficientNets. *Biomed Signal Process Control*. 2023;84: 104777. <https://doi.org/10.1016/j.bspc.2023.104777>.
- [18] Alsaif, H., Guesmi, R., Alshammari, B. M., Hamrouni, T., Guesmi, T., Alzamil, A., et al. (2022). A novel data augmentation-based brain tumor detection using convolutional neural network. *Appl. Sci.* 12, 3773. doi: 10.3390/app12083773
- [19] Tiwari, A.; Srivastava, S.; Pant, M. Brain tumor segmentation and classification from magnetic resonance images: Review of selected methods from 2014 to 2019. *Pattern Recognit. Lett.* 2020, 131, 244–260.
- [20] Cui H, Li Y, Wang Y, Xu D, Wu LM, Xia Y. Towards accurate cardiac MRI segmentation with variational autoencoder-based unsupervised domain adaptation. *IEEE Trans Med Imaging*. 2024. <https://doi.org/10.1109/TMI.2024.3382624>.
- [21] Butt UM, Arif R, Letchmunan S, Malik BH, Butt MA. Feature enhanced stacked auto encoder for diseases detection in brain MRI. *Comput Mater Continua*. 2023. <https://doi.org/10.3260/cmc.2023.039164>
- [22] Ullah F, Nadeem M, Abrar M, Al-Razgan M, Alfakih T, Amin F, Salam A. Brain tumor segmentation from MRI images using handcrafted convolutional neural network. *Diagnostics*. 2023;13(16):2650. <https://doi.org/10.3390/diagnostics13162650>
- [23] Nickparvar, M. (2024). *Brain Tumor MRI Dataset*. Kaggle. <https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset>