

# BRAIN HAEMORRHAGE ANALYSIS USING DEEP LEARNING

Aruna Kokkula<sup>1</sup> and P. Chandra Sekhar<sup>2</sup>

<sup>1</sup>Department of Electronics and Communication Engineering, Maturi Venkata Subba Rao Engineering College, India

<sup>2</sup>Department of Electronics and Communication Engineering, University College of Engineering, Osmania University, India

## Abstract

*Brain haemorrhage remains a critical medical condition with high mortality and disability rates, necessitating timely and precise diagnosis. Traditional diagnostic approaches such as CT imaging often suffer from delays and subjectivity due to the reliance on radiologist interpretation. The main goal of this study is to create a deep learning-driven system that can automatically and reliably identify and categorize brain haemorrhages. The research addresses key challenges such as diagnostic delays, inconsistencies between medical evaluations, and the necessity for scalable and efficient diagnostic methods. This work aims to bridge the existing knowledge gap in real-time, generalized haemorrhage detection across diverse imaging scenarios. This work introduces a unique hybrid deep learning framework that integrates EfficientNetB0 with Bidirectional LSTM and Multi-Head Attention components. The EfficientNetB0 component efficiently extracts spatial features from CT images. These features are reshaped into temporal sequences and processed by BiLSTM to capture bidirectional dependencies. Subsequently, Multi-Head Attention is applied to focus dynamically on significant sequence segments, with residual connections enhancing stability. The training process employs the Adam optimization algorithm along with categorical cross-entropy loss enhanced by label smoothing for improved performance. Training is further regulated through dropout, early stopping, and learning rate scheduling—ensuring robustness. The combination of these elements enhances both the originality and performance of the suggested framework. Experimental results demonstrate that the model attains a classification accuracy of 98.03% and an F1-score of 0.99, surpassing traditional architectures like ResNet50, MobileNet, and DenseNet. Confusion matrix analysis demonstrates minimal false predictions, underscoring high sensitivity and specificity. These findings indicate that the model holds strong potential for use in clinical environments, especially where access to radiological expertise is limited. The integration of convolutional, sequential, and attention-based mechanisms significantly enhances diagnostic performance, offering an intelligent, a scalable approach aimed at enhancing diagnosis and treatment outcomes for individuals with potential brain haemorrhages.*

## Keywords:

*Brain Haemorrhage, Deep Learning, EfficientNetB0, Bidirectional LSTM, Multi-Head Attention, Medical Image Classification*

## 1. INTRODUCTION

Brain haemorrhage is a severe and potentially fatal condition caused by bleeding within the brain due to ruptured blood vessels. It stands among the top causes of death and long-term disability across the globe [1]. Representing a considerable share of stroke-related cases, it necessitates rapid detection and treatment to enhance survival and recovery outcomes. Diagnostic techniques like computed tomography (CT) and magnetic resonance imaging (MRI) are commonly employed to detect and assess haemorrhagic incidents in clinical settings [2]. Nonetheless, these imaging modalities rely heavily on the skill and availability of radiology professionals, which can pose significant challenges, particularly in emergency scenarios or areas with limited medical resources.

With the rapid surge in medical imaging data and the growing complexity of clinical conditions, artificial intelligence (AI) has garnered significant interest for its potential to improve diagnostic precision [3]. Deep learning, in particular, has proven effective in identifying complex patterns within high-dimensional datasets. Its use in analyzing medical images has led to encouraging outcomes in fields like tumor identification, anatomical segmentation, and disease diagnosis [4]. Increasingly, deep learning is being investigated for its ability to assist in the automated detection of brain haemorrhages from neuroimaging scans, aiming to facilitate faster and more informed clinical decisions

Despite these advancements, several challenges continue to hinder the reliable deployment of automated systems in real-world clinical environments. Brain haemorrhages present with diverse appearances, locations, and severities, often complicated by the presence of overlapping pathologies or imaging artifacts [5]. Moreover, the scarcity of large, annotated datasets and the inherent variability in scan quality pose significant obstacles to the development of robust models. Misclassification or delayed detection can have critical implications, underscoring the need for solutions that are both highly accurate and interpretable [6].

Moreover, the evolving characteristics of brain injuries and the requirement to accurately differentiate between various haemorrhage types—such as subarachnoid, subdural, epidural, and intraparenchymal—necessitate advanced analytical methods that surpass the limitations of traditional image processing approaches. The heterogeneity in clinical presentations and patient demographics adds another layer of complexity [7]. These limitations call for the integration of sophisticated learning architectures that can generalize across diverse datasets and adapt to varying clinical requirements with minimal human intervention.

In response to these critical challenges, this research presents a deep learning framework designed to automate the assessment of brain haemorrhages. Utilizing convolutional neural networks (CNNs) trained on labeled neuroimaging data, the model is developed to accurately identify haemorrhage types and detect bleeding [8]. This AI-driven approach not only improves diagnostic accuracy but also minimizes reliance on manual interpretation, making it a scalable and efficient alternative. Comprehensive testing and performance analysis highlight the framework's strong potential to transform early-stage detection and classification of brain haemorrhages in clinical practice.

## 2. LITERATURE SURVEY

Payal Malik et al. [9] emphasized the importance of early and accurate diagnosis in the management of brain haemorrhages. Their study investigated the application of deep learning techniques for identifying haemorrhagic conditions from medical

imaging data. The study evaluated models based on accuracy, sensitivity, specificity, and computational efficiency. Results showed that deep learning significantly aids in diagnosing this life-threatening condition. EfficientNetB3 outperformed other models with 99.95% training and 93.29% validation accuracy. Despite its complexity, EfficientNetB3 remained efficient and well-suited for transfer learning. Other models tested included EfficientNetB2, ResNet, SEResNext, and ResNext, with lower performance.

Shanu Nizarudeen et al. [10] investigated deep learning-based approaches for identifying intracranial haemorrhage (ICH), a condition requiring rapid medical attention. Their research involved evaluating models such as HResNet, HResNet-SE, and HRaNet using 22,811 CT scans from the CQ500 dataset. Among these, HRaNet achieved the best results, recording an AUC above 0.96, a Jaccard index of 0.9130, and Macro and Micro F1-scores of 0.9464 and 0.9545, respectively. It also achieved a Kappa score over 0.9063, indicating strong reliability. The study tackled multi-label classification for identifying multiple haemorrhage types. It highlighted HRaNet's efficiency in both accuracy and prediction time. These findings show deep learning's promise in aiding ICH diagnosis and improving clinical outcomes.

Cansu Yalcin et al. [11] introduced a deep learning-based method aimed at forecasting hematoma expansion (HE) in patients with intracerebral haemorrhage (ICH) by analysing their initial non-contrast CT images. HE, seen in 30–38% of cases within 24 hours, worsens patient outcomes but is difficult to predict due to limited data. The study used a 2D Efficient Net B0 model trained on 122 patient scans, including 35 HE cases, with manual lesion annotations. To address data scarcity, synthetic images were added to the training set. The best results came from using five synthetic versions per image with standard augmentation. This boosted accuracy from 0.56 to 0.84 and F1-score from 0.53 to 0.82. The method shows strong potential to aid early HE prediction and improve clinical management.

Anandakumar Haldorai et al. [12] highlighted the potential of deep learning in medical diagnostics, particularly for extracting critical insights from healthcare data. The study underscored the need for timely evaluation of CT scans by radiologists to facilitate early identification of cerebral haemorrhages. Utilizing a dataset containing full-body DICOM CT scans from 3,000 individuals, brain images were isolated using segmentation techniques. These segmented scans were grouped based on visual features to improve both the precision and speed of haemorrhage detection. This clustering strategy contributed to more consistent diagnostic outcomes. Furthermore, a convolutional neural network (CNN) was applied to classify the brain CT images, effectively distinguishing between haemorrhagic and non-haemorrhagic cases.

Andrea Zirn et al. [13] investigated the application of post-mortem CT (PMCT) imaging for identifying fatal cerebral haemorrhages, emphasizing its potential as a rapid and non-invasive complement to traditional autopsy methods. The study analysed 81 PMCT cases, including 36 confirmed haemorrhage cases and 45 neurologically normal controls, to train and validate a set of six machine learning algorithms alongside two deep learning architectures—CNN and DenseNet. An 80/20 train-validation split with five-fold cross-validation was used to assess model performance. The convolutional neural network (CNN)

outperformed all others, achieving an accuracy of 0.94. These findings highlight the capability of deep learning models in enhancing post-mortem diagnostic processes and suggest that CNNs can support forensic pathologists in determining causes of death using PMCT scans, promoting more efficient and automated autopsy practices.

Chi-Tung Cheng et al. [14] reviewed the growing role of deep learning (DL) in trauma imaging, highlighting its effectiveness in detecting injuries across modalities like FAST, X-rays, and CT scans. DL has been applied to identify intracranial haemorrhages, vertebral fractures, and organ damage in trauma patients. The study outlines core principles for DL algorithm development and showcases current clinical applications. Future directions include federated learning for diverse datasets, improved model transparency, and integration of multimodal data for better insights. Although some FDA-approved AI tools exist, clinical adoption remains limited. The study emphasizes the importance of cross-disciplinary cooperation in developing effective and clinically validated tools. Overall, deep learning demonstrates significant promise in enhancing trauma diagnosis and improving patient care outcomes.

A.S. Neethi et al. [15] investigated deep learning approaches for the automatic detection of haemorrhagic stroke, highlighting the critical need for timely and precise diagnosis. Their research involved a systematic review and gap analysis, leveraging extensive publicly available datasets of non-contrast brain CT images. By evaluating across multisite data and varying haemorrhage complexities, key challenges and limitations in current models were identified. Results revealed gaps in accuracy and reliability, underscoring the need for enhanced techniques. The research highlights how AI can support radiologists and improve diagnostic efficiency. It also offers guidance for future development of robust diagnostic tools. Overall, the study aims to advance sustainable and effective stroke care.

Süleyman Uzun et al. [16] explored AI-based detection of brain strokes, focusing on real-time analysis of CT scans using YOLO models. Strokes, particularly common in those over 65, are classified as ischemic or haemorrhagic, with early diagnosis being critical. The study compared YOLOv7, YOLOv8, and YOLOv9 against U-Net and Mask-RCNN using 6,951 anonymised CT slices. YOLOv9-Seg achieved the highest performance, with mAP@0.5 scores of 99.50% for ischemic, 99.49% for haemorrhagic, and 99.71% for combined stroke cases. Models were trained using Pytorch and CUDA acceleration. Results show YOLOv9-Seg's superiority in accuracy and speed, making it ideal for emergency stroke detection. The findings highlight the model's potential to support rapid, AI-assisted diagnosis in clinical settings.

### 3. PROPOSED MODEL

The developed model introduces a hybrid deep learning framework that integrates EfficientNetB0, Bidirectional LSTM, and Multi-Head Attention to improve the accuracy of brain haemorrhage classification. Initially, grayscale medical images are converted into three-channel inputs to align with the requirements of EfficientNetB0, a high-performance pretrained CNN known for extracting detailed spatial features. These spatial features are then reshaped into a sequential format, enabling a

Bidirectional LSTM to learn contextual relationships in both temporal directions. To enhance feature relevance, a Multi-Head Attention module is employed, allowing the model to concentrate selectively on crucial parts of the sequence while maintaining contextual integrity via residual connections. The attention-weighted outputs are subsequently flattened and fed through fully connected layers, with dropout applied to mitigate overfitting, leading to a softmax layer for binary classification. Training optimization is carried out using the Adam algorithm, complemented by label smoothing for better generalization. Additionally, the model employs early stopping and adaptive learning rate scheduling to achieve stable and efficient convergence.

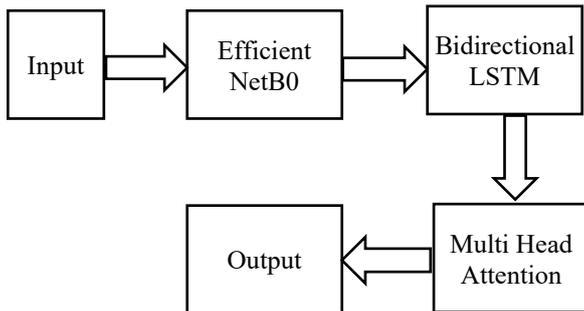


Fig.1, Structural Overview of the Proposed Hybrid Framework for Brain Haemorrhage

The Fig.1 presents the structure of the proposed hybrid deep learning model tailored for detecting brain haemorrhages. The workflow starts with input medical images that undergo preprocessing before being passed into EfficientNetB0—a robust CNN pretrained on ImageNet—to extract meaningful spatial features. These features are then reformatted into a sequential structure and processed through a Bidirectional LSTM layer, which captures bidirectional temporal relationships to deepen contextual interpretation. To further refine the learned features, a Multi-Head Attention module is employed, enabling the model to selectively emphasize the most relevant information across multiple representation subspaces. The refined features are finally passed to a dense output layer for classification, enabling accurate prediction of brain haemorrhage presence and type. This hybrid architecture effectively combines convolutional, recurrent, and attention-based mechanisms to achieve robust diagnostic performance.

### 3.1 EFFICIENTNET-B0

EfficientNetB0, developed by Google AI, serves as the foundational model in the EfficientNet series, known for its balance between performance and computational efficiency. It employs a unique compound scaling strategy that simultaneously adjusts the network's depth, width, and input resolution to optimize accuracy and resource usage. The architecture is based on Mobile Inverted Bottleneck Convolutions (MBCConv) combined with squeeze-and-excitation blocks, which contribute to its lightweight design and reduced computational cost. Although it is the most compact model in the Efficient Net family, EfficientNetB0 surpasses larger architectures like ResNet-50 in both classification accuracy and inference speed, making it particularly suitable for applications on mobile and edge devices with limited processing power.

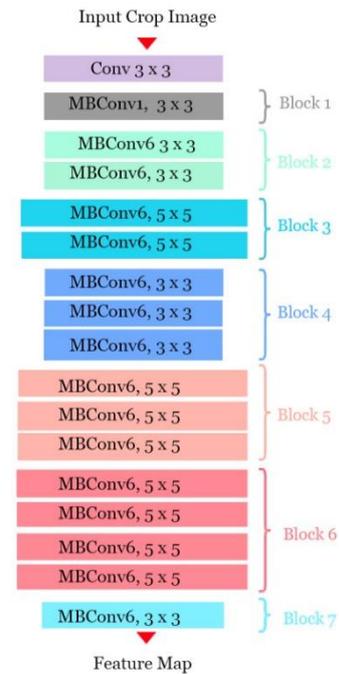


Fig.2. Architecture for EfficientNet-B0

The Fig.2 showcases a deep convolutional neural network architecture inspired by EfficientNet, tailored to extract detailed feature representations from crop images. The network is structured into seven consecutive blocks, comprising standard convolutional and Mobile Inverted Bottleneck Convolution (MBCConv) layers, with kernel sizes of  $3 \times 3$  and  $5 \times 5$ . The model begins with a conventional convolution layer, followed by MBCConv layers distributed across Blocks 1 to 7. Among these, MBCConv6 layers are prominently featured to enhance model depth and computational efficiency, utilizing depthwise separable convolutions to lower processing demands. As the input passes through each block, the network captures progressively richer spatial and semantic features. The resulting "Feature Map" provides a high-level abstraction of the original image, making it suitable for advanced tasks such as classification, object detection, or image segmentation.

### 3.2 BIDIRECTIONAL LSTM

A Bidirectional Long Short-Term Memory (BiLSTM) network extends the conventional LSTM by processing sequence data in two directions—forward and backward—thereby offering a more comprehensive understanding of contextual information. Unlike standard LSTMs, which operate solely from past to future, BiLSTM consists of two separate LSTM layers: one handling the sequence in its original order, and the other processing it in reverse. This structure enables the model to incorporate insights from both preceding and succeeding elements at each time step. Such an approach is highly effective in applications like language modeling, speech processing, and biomedical signal interpretation, where capturing temporal dependencies from both directions enhances predictive accuracy.

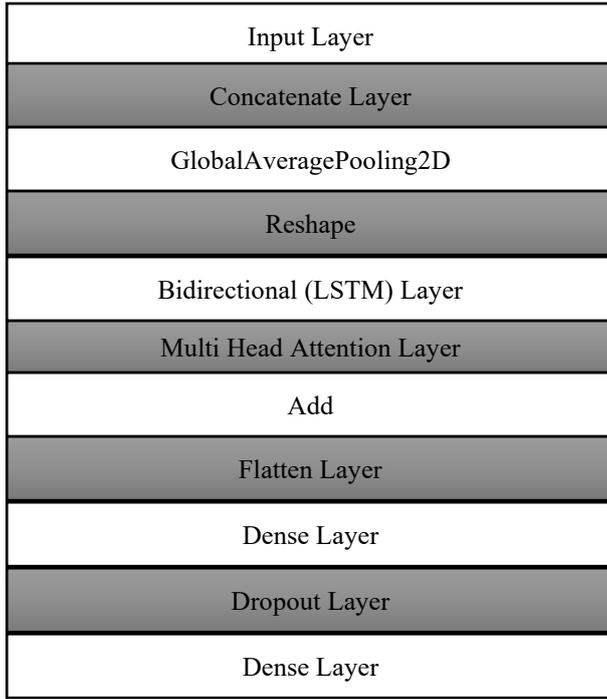


Fig.3. Architecture of the Proposed Approach

### 3.2.1 Input Layer (Input):

This layer specifies the dimensions of the incoming data. Here, the model is configured to accept grayscale images with a resolution of 256 by 256 pixels, represented by the input shape (256, 256, 1). It acts as the entry point into the neural network and does not perform any computation itself. This layer ensures the model knows what type of data to expect during training and inference.

$$X_{\text{rgb}} = \text{Concatenate}(X, X, X) \in \mathbb{R}^{256 \times 256 \times 3}$$

### 3.2.2 Concatenate Layer (Concatenate):

The Concatenate layer is used to combine multiple tensors along a specified axis. Here, it duplicates the single-channel grayscale image three times to create a 3-channel tensor. This is necessary because EfficientNetB0, which is pretrained on RGB images (3 channels), expects input data with three channels. The result is a (256, 256, 3) image tensor.

### 3.2.3 Global Average Pooling Layer (Global Average Pooling2D):

GlobalAveragePooling2D simplifies each feature map by computing the average of all its elements, resulting in a single value per channel. Positioned after convolutional layers, it transforms the 3D feature output (height  $\times$  width  $\times$  channels) into a 1D vector consisting only of channel values. This operation helps lower model complexity and minimizes the risk of overfitting. It retains spatially summarized global features and is often used before fully connected layers.

$$v = \frac{1}{H \cdot W} \sum_{i=1}^H \sum_{j=1}^W F_{i,j} \in \mathbb{R}^{1280}$$

### 3.2.4 Reshape Layer (Reshape):

The Reshape layer changes the shape of its input tensor without altering its data. In your model, it reshapes the 1280-

dimensional vector (from Global Average Pooling) into a shape of (10, 128), simulating 10 time steps with 128 features each. This transformation prepares the data for processing by the LSTM, which requires sequential input.

$$S = \text{Reshape}(v) \in \mathbb{R}^{10 \times 128}$$

### 3.2.5 Bidirectional LSTM Layer (Bidirectional (LSTM)):

The Bidirectional wrapper allows an LSTM layer to analyse sequence data in two directions—forward and backward—capturing contextual information from both the past and future. It internally runs two separate LSTM layers on the input sequence (one in the original order and one in reverse) and merges their outputs. This dual processing enhances the model's ability to recognize temporal patterns within the transformed feature sequences.

- The BiLSTM layer captures context by traversing the sequence in both directions:

$$\vec{h}_t = \text{LSTM}_f(s_t)$$

$$\overleftarrow{h}_t = \text{LSTM}_b(s_t)$$

$$h_t = [\vec{h}_t; \overleftarrow{h}_t] \in \mathbb{R}^{10 \times 128}$$

- Output sequence:

$$H = [h_1, h_2, \dots, h_{10}] \in \mathbb{R}^{10 \times 128}$$

### 3.2.6 Multi-Head Attention Layer (Multi Head Attention):

This layer employs the attention mechanism across multiple parallel heads (in this case, two), with each head learning to attend to different segments of the input sequence. This allows the model to identify diverse contextual relationships within the data. By assigning dynamic importance to various time steps or features, attention mechanisms enhance the model's performance in handling sequential tasks such as this one.

- Let  $Q=K=V=H$  (self-attention):

$$\text{Attention}(Q, K, V) = \text{softmax} \left( \frac{QK^T}{\sqrt{d_k}} \right) V$$

- Multi-head attention with 2 heads:

$$\text{MHA}(H) = \text{Concat}(\text{head}_1, \text{head}_2) W^O \in \mathbb{R}^{10 \times 128}$$

- Residual connection:

$$A = H + \text{MHA}(H)$$

### 3.2.7 Add Layer (Add):

The Add layer performs element-wise addition of two tensors. In this model, it implements a residual connection by adding the original LSTM output with the output of the MultiHeadAttention. This helps in stabilizing training and preserving the original sequential features while integrating the learned attention patterns, much like in Transformer or ResNet architectures.

### 3.2.8 Flatten Layer (Flatten):

The Flatten layer transforms a multi-dimensional tensor into a one-dimensional array, preparing the data for input into fully connected dense layers. Following the temporal feature extraction, it compresses the entire sequence structure into one long vector, enabling classification or regression tasks using dense layers.

Flatten the sequence:

$$A_f = \text{Flatten}(A) \in \mathbb{R}^{1280}$$

### 3.2.9 Dense Layer (Dense(256)):

This dense layer, consisting of 256 units and using ReLU activation, is designed to learn intricate patterns from the input features. By connecting each neuron to all elements of the flattened input, the layer enables high-level interpretation of the integrated spatial and temporal information extracted by the preceding layers.

Dense layer with ReLU:

$$z_1 = \text{ReLU}(W_1 A_f + b_1) \in \mathbb{R}^{256}$$

### 3.2.10 Dropout Layer (Dropout(0.4)):

During training, the Dropout layer deactivates 40% of its input units at random. This regularization strategy reduces overfitting by preventing the model from relying too heavily on particular neurons. As a result, it encourages the learning of more generalizable and resilient features that perform well on new, unseen data.

Dropout:

$$z_1' = \text{Dropout}(z_1, \text{rate} = 0.4)$$

### 3.2.11 Output Dense Layer (Dense(2, softmax)):

The final Dense layer contains two output neurons, each representing one of the target classes. Using the softmax activation function, it transforms the outputs into probability scores for each class. The class with the highest probability is selected as the model's predicted label. This configuration is well-suited for binary classification tasks.

Output layer (Softmax):

$$\hat{y} = \text{softmax}(W_2 z_1' + b_2) \in \mathbb{R}^2$$

## 3.3 OPTIMIZATION FUNCTION (ADAM OPTIMIZER)

The model employs the Adam optimizer (Adaptive Moment Estimation) to adjust its weights throughout the training process. Adam is a refined version of stochastic gradient descent that assigns unique learning rates to each parameter by adaptively estimating the first and second moments of the gradients. At each iteration  $t$ , it computes the exponentially weighted moving averages of the gradients ( $m_t$  for the mean) and their squares ( $v_t$  for the variance), followed by bias correction. The parameters are then updated based on these corrected estimates using the following rule:

$$\begin{aligned} m_t &= \beta_1 m_{t-1} + (1 - \beta_1) g_t \\ v_t &= \beta_2 v_{t-1} + (1 - \beta_2) g_t^2 \\ \hat{m}_t &= \frac{m_t}{1 - \beta_1^t}, \quad \hat{v}_t = \frac{v_t}{1 - \beta_2^t} \\ \theta_t &= \theta_{t-1} - \eta \cdot \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \delta} \end{aligned}$$

In this context,  $\theta$  denotes the model's parameters, while  $g_t$  is the computed gradient at iteration  $t$ . The learning rate is represented by  $\eta$ , and  $\beta_1$  and  $\beta_2$  are decay rate hyperparameters for

the exponential moving averages—commonly set to  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . A small constant  $\epsilon$  (typically  $10^{-8}$ ) is added to ensure numerical stability. For this model, the learning rate is set to a constant value of  $1 \times 10^{-4}$ .

## 3.4 LOSS FUNCTION (CATEGORICAL CROSS ENTROPY WITH LABEL SMOOTHING)

The training process utilizes categorical cross-entropy as the loss function, incorporating label smoothing to enhance generalization. This method softens the target labels by replacing strict 0s and 1s with slightly adjusted values, which discourages the model from making overconfident predictions and helps improve robustness. Instead of using target labels like  $[1, 0]$ , it uses softened targets like  $[0.9, 0.1]$ . The smoothed label  $y_i^{\text{smooth}}$  is calculated as:

$$y_i^{\text{smooth}} = (1 - \delta) y_i + \frac{\delta}{C}$$

where,  $\epsilon$  denotes the smoothing coefficient (set to 0.1 in this instance),  $C$  represents the total number of classes (which is 2), and  $y_i$  corresponds to the original one-hot encoded label. With label smoothing applied, the categorical cross-entropy loss is modified as follows:

$$L = - \sum_{i=1}^C y_i^{\text{smooth}} \cdot \log(\hat{y}_i)$$

This loss function improves generalization and helps reduce overfitting, particularly in cases where training data is limited or noisy. The soft targets encourage the model to output probability distributions that are less extreme, thus improving calibration and robustness.

The proposed model incorporates various advanced deep learning techniques for efficient image processing and classification. It begins by accepting a grayscale image of size (256, 256, 1), which is then duplicated across three channels to match the input format required by the pre-trained EfficientNetB0 feature extractor. EfficientNetB0, known for its computational efficiency, serves as the backbone for feature extraction, capturing high-level spatial features from the image without the final classification layer. This setup enables the model to leverage rich image features. To consolidate these features, Global Average Pooling is applied, reducing the feature map to a single vector, making it suitable for sequential analysis. The reshaped output forms a sequence with 10-time steps, each containing 128 features, preparing the data for temporal modelling.

For sequence processing, a Bidirectional LSTM layer is employed, allowing the model to capture both past and future dependencies in the data. To enhance its focus on the most relevant features at each time step, a Multi-Head Attention mechanism is used, followed by a residual connection that combines the attention output with the original input, preserving important features. Finally, the model utilizes several fully connected layers with dropout regularization to combat overfitting, and the output layer employs softmax activation for multi-class classification. The model is optimized using the Adam optimizer with a low learning rate and categorical crossentropy loss with label smoothing for stable learning. To further enhance training, early stopping and dynamic learning rate reduction

callbacks are applied to prevent overfitting and adapt the learning rate based on validation performance.

### 4. EXPERIMENTAL RESULTS

In this subsection, we present an extensive examination of the results from the proposed approach for the ongoing simulations. The dataset used for these simulations was obtained from the Brain CT Haemorrhage Dataset [17]. The data processing techniques outlined above were applied to the dataset used for this analysis in the study

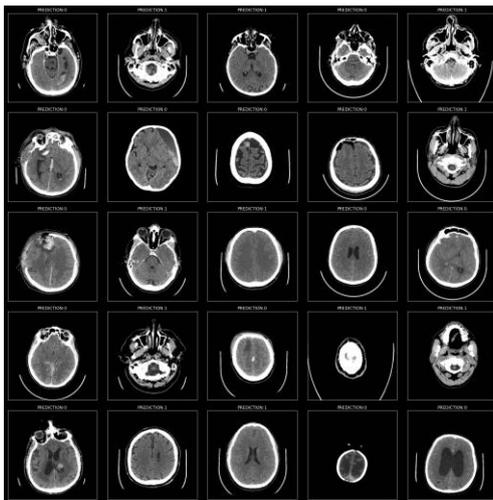


Fig.4. Deep Learning-Based Classification of Brain Haemorrhage from CT Images

This Fig.4 displays a series of brain CT scan slices arranged in a grid format, each annotated with the output of a deep learning model designed to classify brain haemorrhages. Every image tile represents a unique axial view, labelled as either 'PREDICTION: likely normal' or 'Prediction: indicative of haemorrhage.' The model differentiates between normal and abnormal regions by analyzing variations in tissue density and structural irregularities. Such AI-driven classification tools play a vital role in emergency radiology, supporting rapid patient triage and assisting healthcare professionals in the early detection and treatment of intracranial bleeding.

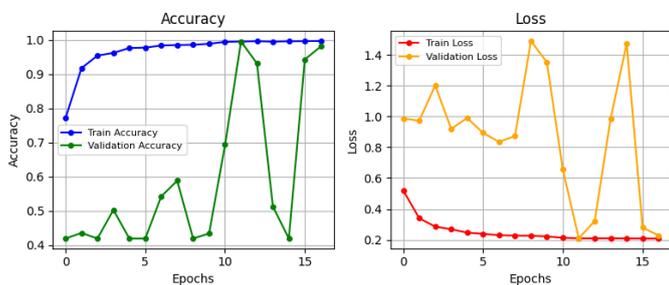


Fig.5. Epoch-wise Training and Validation Performance Metrics

The Fig.5 illustrates the deep learning model's training and validation performance across 16 epochs, highlighting both accuracy and loss trends. In the left plot, training accuracy steadily improves, approaching near-perfect levels, whereas validation accuracy shows noticeable fluctuations—pointing to potential overfitting or instability. The right plot reveals a smooth

decline in training loss, indicating effective learning, while the validation loss varies significantly, and reinforcing concerns about generalization. The gap between training and validation results suggests a need for regularization strategies such as dropout, data augmentation, or early stopping to improve the model's robustness.

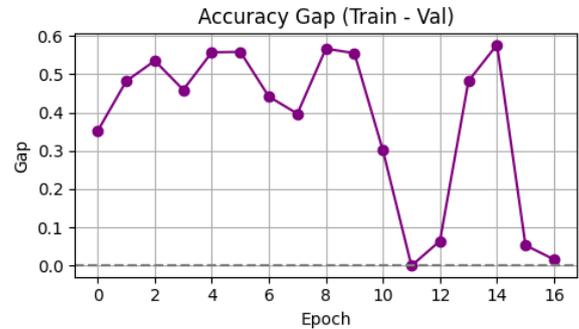


Fig.6. Epoch-wise Comparison of Training and Validation Accuracy

The Fig.6 displays the difference in accuracy between the training and validation stages over 17 epochs. This accuracy gap is calculated by subtracting the validation accuracy from the training accuracy at each epoch and serves as a measure of potential overfitting. The consistently large differences—often exceeding 0.4—indicate that the model is significantly overperforming on the training data compared to validation. Although the gap momentarily narrows during epochs 11, 12, 15, and 16, the overall pattern of sharp variations suggests unstable generalization. These observations highlight the need for improved regularization methods, refined validation practices, or the implementation of cross-validation techniques to promote better model generalization on unseen data.

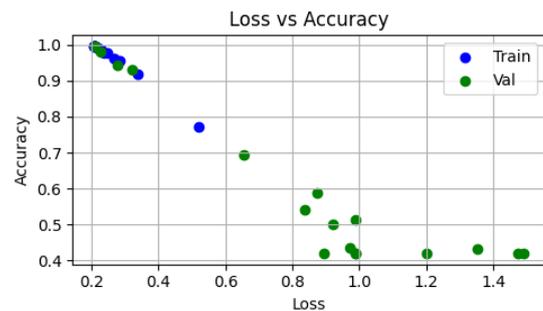


Fig.7. Scatter Plot of Loss vs Accuracy for Training and Validation

The Fig.7 displays a scatter plot illustrating the correlation between loss and accuracy for both training and validation datasets across epochs. Each point corresponds to the model's performance at a specific epoch. The blue markers, representing training data, reveal a clear inverse trend—lower loss values are associated with higher accuracy, reflecting effective learning. In contrast, the green markers, corresponding to the validation set, exhibit a more irregular distribution, with no consistent pattern and several instances of high loss coupled with low accuracy. This irregularity indicates the model's difficulty in generalizing beyond the training data, pointing to potential overfitting. The visible separation between training and validation points further

emphasizes the performance gap previously observed in earlier evaluations.

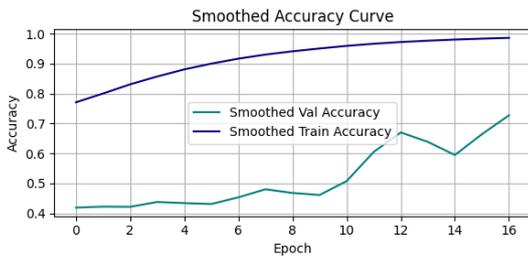


Fig.8. Smoothed Accuracy Curve across Epochs

The Fig.8 illustrates smoothed accuracy trends for both training and validation sets across 17 epochs. A smoothing technique is applied to minimize fluctuations and better reveal overall learning behavior. The dark blue line, representing the training accuracy, shows a consistent upward trend, indicating effective learning and improved model performance on the training data. Meanwhile, the teal line for validation accuracy shows a more gradual and erratic progression, with occasional declines pointing to difficulties in generalization. The increasing separation between the two curves signals overfitting, where the model adapts too closely to the training data while underperforming on new examples. This pattern underscores the necessity of strategies like early stopping and regularization to enhance generalization.

Table.1. Classification Report

	Precision	Recall	F1-Score
Haemorrhage	1.00	0.98	0.99
Normal	0.98	1.00	0.99
Accuracy	0.99		

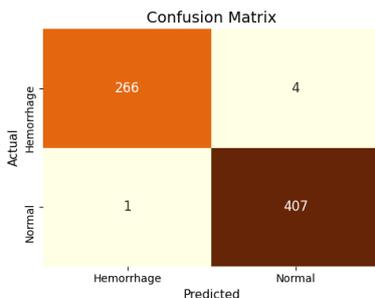


Fig.9. Confusion Matrix

The Table.1 summarizes the model’s classification performance in distinguishing between brain haemorrhage and normal cases. The results indicate outstanding precision and recall across both classes, with precision scores of 1.00 for haemorrhage and 0.98 for normal, and recall values of 0.98 and 1.00, respectively. These metrics yield an F1-score of 0.99 for each class, reflecting an excellent balance between precision and recall. The model achieves an overall accuracy of 99%, highlighting its effectiveness in accurately classifying most input images with minimal errors. These performance indicators suggest the model is highly dependable and well-suited for deployment in sensitive medical diagnostic scenarios.

The Fig.9 displays the confusion matrix for a binary classification model tasked with differentiating between Haemorrhage and Normal cases. According to the matrix, the model correctly classified 266 out of 270 Haemorrhage cases, with only 4 incorrectly labeled as Normal. For the Normal class, 407 out of 408 instances were accurately identified, with just one false positive. These results highlight the model’s excellent classification ability, marked by very few misclassifications and demonstrating both high sensitivity and specificity. The prominent diagonal values confirm the model’s reliability and effectiveness in accurately distinguishing brain haemorrhage cases from normal scans.

Table.2. Comparative Analysis

Methods	Accuracy
VGG16 [17]	70%
ResNet50 [18]	81.58%
AlexNet [19]	84.33%
DenseNet [20]	88.01%
VGG19 [21]	91.56%
MobileNet [22]	93.79%
Proposed Model (hybridmodel)	98.03

The Table.2 provides a comparison of different deep learning models evaluated by their classification accuracy. Among the evaluated baseline models, VGG16 and ResNet50 achieved relatively lower accuracies of 70% and 81.58% respectively, while AlexNet and DenseNet performed better with accuracies of 84.33% and 88.01%. VGG19 and MobileNet demonstrated further improvement, reaching 91.56% and 93.79% respectively. Notably, the proposed hybrid model outperformed all existing architectures with a significant accuracy of 98.03%, indicating its superior feature extraction and classification capabilities. This highlights the effectiveness of integrating CNN-based feature extraction with sequence modeling components, offering a robust solution for the target task.

## 5. CONCLUSION

This research introduces and validates an innovative hybrid deep learning framework aimed at the automated detection and classification of brain haemorrhages using CT scan data. The model integrates EfficientNetB0 for extracting high-level spatial features, a Bidirectional Long Short-Term Memory (BiLSTM) layer for modeling temporal dependencies, and a Multi-Head Attention mechanism to emphasize the most relevant feature sequences. This combination equips the network to effectively recognize complex haemorrhagic patterns, which may vary in shape, intensity, and anatomical location. To enhance model generalization and stability, the architecture incorporates residual connections, dropout for regularization, label smoothing, and adaptive learning techniques such as the Adam optimizer and early stopping. Evaluation results show the model achieving a high classification accuracy of 98.03% and an F1-score of 0.99, surpassing several established architectures including VGG16, ResNet50, and MobileNet. The confusion matrix confirms minimal false predictions, demonstrating strong performance in distinguishing haemorrhagic from normal cases. These findings

highlight the model's potential as a precise, interpretable, and scalable solution for clinical applications—particularly in urgent or low-resource environments. By integrating convolutional, recurrent, and attention-based learning strategies, this study makes a notable contribution to AI-driven neuroimaging and paves the way for further advancements in automated medical diagnosis systems.

## REFERENCES

- [1] Matteo Antonio Sacco, Maria Cristina Verrina, Roberto Raffaele, Saverio Gualtieri, Alessandro Pasquale Tarallo, Santo Gratteri and Isabella Aquila, "The Role of Autopsy in the Forensic and Clinical Evaluation of Head Trauma and Traumatic Brain Injury in Road Traffic Accidents: A Review of the Literature", *Diagnostics*, Vol. 15, No. 4, pp. 1-15, 2025.
- [2] Nizarudeen Shanu and Ganesh Ramaswamy Shanmughavel, "Comparative Analysis of ResNet, ResNet-SE and Attention-based RaNet for Haemorrhage Classification in CT Images using Deep Learning", *Biomedical Signal Processing and Control*, Vol. 88, pp. 1-7, 2024.
- [3] Mohamed Khalifa and Mona Albadawy, "AI in Diagnostic Imaging: Revolutionising Accuracy and Efficiency", *Computer Methods and Programs in Biomedicine Update*, Vol. 5, pp. 1-12, 2024.
- [4] Yan Xu, Rixiang Quan, Weiting Xu, Yi Huang, Xiaolong Chen and Fengyuan Liu, "Advances in Medical Image Segmentation: A Comprehensive Review of Traditional, Deep Learning and Hybrid Approaches", *Bioengineering*, Vol. 11, No. 10, pp. 1-42, 2024.
- [5] Anshu Mahajan and Ashima Mahajan, "Neuroimaging: CT Scan and MRI", *Principles and Practice of Neurocritical Care*, pp. 189-215, 2024.
- [6] A.A. Talal Abdullah, Mohd Soperi Mohd Zahid and Waleed Ali, "A Review of Interpretable ML in Healthcare: Taxonomy, Applications, Challenges and Future Directions", *Symmetry*, Vol. 13, No. 12, pp. 1-9, 2021.
- [7] A. Marcel Heinrich, M.R.H. Ahmed Mostafa, P. Jennifer Morton, J.A.C. Lukas Hawinkels and Jai Prakash, "Translating Complexity and Heterogeneity of Pancreatic Tumor: 3D in Vitro to in Vivo Models", *Advanced Drug Delivery Reviews*, Vol. 174, pp. 265-293, 2021.
- [8] Quoc Tuan Hoang, Xuan Hien Pham, Xuan Thang Trinh, Anh Vu Le, V. Minh Bui and Trung Thanh Bui, "An Efficient CNN-based Method for Intracranial Haemorrhage Segmentation from Computerized Tomography Imaging", *Journal of Imaging*, Vol. 10, No. 4, pp. 1-15, 2024.
- [9] Payal Malik, Ajay Dureja, Aman Dureja, Rajkumar Singh Rathore and Nisha Malhotra, "Enhancing Intracranial Haemorrhage Diagnosis through Deep Learning Models", *Procedia Computer Science*, Vol. 235, pp. 1664-1673, 2024.
- [10] Cansu Yalcin, Valeriia Abramova, Mikel Terceno, Arnau Oliver, Yolanda Silva and Xavier Llado, "Hematoma Expansion Prediction in Intracerebral Haemorrhage Patients by using Synthesized CT Images in an End-to-End Deep Learning Framework", *Computerized Medical Imaging and Graphics*, Vol. 117, pp. 1-9, 2024.
- [11] Anandakumar Haldorai, Suriya Murugan and Minu Balakrishnan, "Haemorrhage Detection from Whole-Body CT Images using Deep Learning", *Artificial Intelligence for Sustainable Development*, pp. 139-151, 2024.
- [12] Andrea Zirn, Eva Scheurer and Claudia Lenz, "Automated Detection of Fatal Cerebral Haemorrhage in Postmortem CT Data", *International Journal of Legal Medicine*, Vol. 138, No. 4, pp. 1391-1399, 2024.
- [13] Chi-Tung Cheng, Chun-Hsiang Ooyang, Chien-Hung Liao and Shih-Ching Kang, "Applications of Deep Learning in Trauma Radiology: A Narrative Review", *Biomedical Journal*, Vol. 48, No. 1, pp. 1-7, 2025.
- [14] A.S. Neethi, Santhosh Kumar Kannath, Adarsh Anil Kumar, Jimson Mathew and Jeny Rajan, "A Comprehensive Review and Experimental Comparison of Deep Learning Methods for Automated Haemorrhage Detection", *Engineering Applications of Artificial Intelligence*, Vol. 133, pp. 1-9, 2024.
- [15] Suleyman Uzun and Mehmet Okuyar, "A New Deep Learning-based GUI Design and Implementation for Automatic Detection of Brain Strokes with CT Images", *The European Physical Journal Special Topics*, Vol. 234, No. 1, pp. 141-164, 2025.
- [16] A. Abdulkader, "Brain CT Haemorrhage Dataset", Available at <https://www.kaggle.com/datasets/abdulkader90/brain-ct-haemorrhage-dataset>, Accessed in 2025.
- [17] Hamza Sekkat, Abdellah Khallouqi, Omar El Rhazouani and Abdellah Halimi, "Automated Detection of Hydrocephalus in Pediatric Head Computed Tomography using VGG 16 CNN Deep Learning Architecture and based Automated Segmentation Workflow for Ventricular Volume Estimation", *Journal of Imaging Informatics in Medicine*, pp. 1-18, 2025.
- [18] Omid Mirzaei, Sedra Aliasgher Mohammed, Boran Sekeroglu and Ahmet Ilhan, "Comparison of Intracranial Haemorrhages Detection Performances of Deep Learning Models on CT Images", *Procedia Computer Science*, Vol. 258, pp. 3194-3202, 2025.
- [19] Simarjeet Kaur and Amar Singh, "A New Deep Learning Framework for Accurate Intracranial Brain Haemorrhage Detection and Classification using Real-Time Collected NCCT Images", *Applied Magnetic Resonance*, Vol. 55, No. 6, pp. 629-661, 2024.
- [20] Yurui Hu, Tianyu Liu, Shutong Pang, Xiao Ling, Zhanqiu Wang and Wenfei Li, "Deep Learning-Assisted Diagnosis of Placenta Accreta Spectrum using the DenseNet-121 Model: A Multicenter, Retrospective Study", *Journal of Imaging Informatics in Medicine*, pp. 1-10, 2025.
- [21] Mohemmed Sha, "A Graph Neural Network Technique for the Prediction of Cerebral Stroke using an Unbalanced Medical Dataset", *Multimedia Tools and Applications*, pp. 1-37, 2025.
- [22] Yu-Ruei Chen, Chih-Chieh Chen, Chang-Fu Kuo and Ching-Heng Lin, "An Efficient Deep Neural Network for Automatic Classification of Acute Intracranial Haemorrhages in Brain CT Scans", *Computers in Biology and Medicine*, Vol. 176, pp. 1-11, 2024.