International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 09 | Sept - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

Deep Learning Approaches for Multilabel Classification of Brain Hemorrhages in CT Imaging

Veena B[1], Anusha Daivajnya[2]
[1]Assistant Proffesor, Department of MCA, UBDTCE, Davanagere, Karnataka ,India
[2]Student of UBDT College of Engineering, Davanagere, Karnataka,India

ABSTRACT

As a serious neurological emergency, intracranial hemorrhage requires an early and precise diagnosis to improve the patient's chances of recovery. Using the RSNA Intracranial Hemorrhage Challenge dataset, this study introduces a deep learning system for the automated detection and categorization of brain hemorrhages from computed tomography (CT) scans. To improve the data and strengthen the model, preprocessing methods such as contrast improvement, scaling, augmentation, and intensity normalization were used. There were five different kinds of hemorrhages: subarachnoid, intraventricular, intraparenchymal, and epidural. To improve the generalization abilities and lessen the overfitting tendencies of ResNet and DenseNet121, two sophisticated convolutional neural network architectures, we used regularization techniques like dropout, label smoothing, and data augmentation. The results of the experiment showed that both models performed well across all classifications and correctly differentiated between different types of bleeding. The findings imply that decision-making in radiology workflows can be aided by deep learning algorithms. By detecting cerebral hemorrhage more rapidly and precisely, this could result in a quicker diagnosis, less disagreement among observers, and better emergency care.

Keywords: Intracranial Hemorrhage, Deep Learning, CT Imaging, ResNet, DenseNet121, Medical Image Classification.

INTRODUCTION

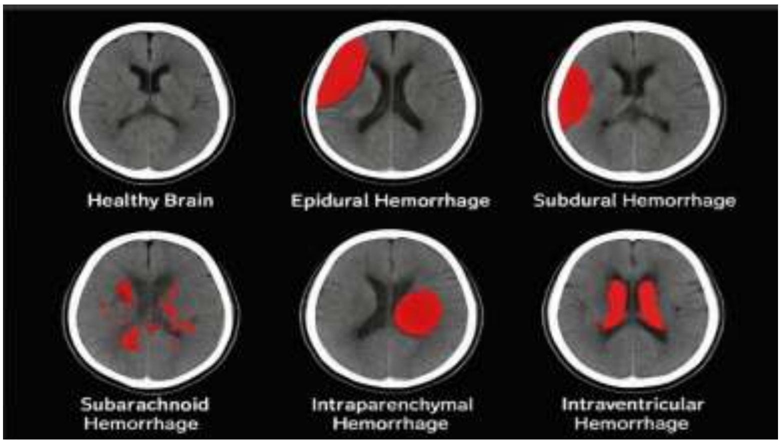

Figure 1: Brain Hemorrhage Subtypes with healthy Brain

An extremely serious medical emergency affecting the brain is intracranial hemorrhage (ICH). Cerebral blood artery rupture within brain tissue or nearby regions causes this disorder, which results in blood accumulation that disrupts normal neurological function. The location, volume, and rate of bleeding propagation all affect the severity of the clinical manifestations, which range from severe neurological impairments to fatal outcomes. According to recent global health data, ICH accounts for over 3.3 million annual deaths worldwide, or 45.6% of all stroke-related deaths. In emergency medical situations, this substantial illness burden necessitates quick diagnostic response and precise subtype classification.
Due to its speed, accessibility in clinics, and superior ability to visualize intracranial structures, computed tomography (CT) imaging is the most crucial test for diagnosing ICH. Since CT scans can determine whether bleeding is present and what kind of bleeding it is within minutes of a patient arriving, they are the primary diagnostic tool used in emergency rooms. However, interpreting CT scans necessitates specialized radiological knowledge, which may cause issues with diagnosis. Diagnostic accuracy may be compromised by human factors like interpretation variability, clinical fatigue, and differences in experience, particularly when evaluating slight hemorrhagic changes or distinguishing between morphologically identical hemorrhage subtypes. Research

efforts to develop automated diagnostic support systems that increase the speed and precision of radiological interpretation have been sparked by these clinical concerns.

Medical image analysis techniques have changed as a result of recent developments in deep learning. While manually designed feature extraction is the foundation of traditional computer-aided diagnosis techniques, deep learning systems use multi-layered neural networks to autonomously learn hierarchical data representations. Significant performance improvements have resulted from this technological advancement in a number of medical imaging applications, including pulmonary pathology screening, ophthalmological disease classification, and oncological tumor detection. For applications involving intracranial bleeding, deep learning models show great promise. They enable accurate subtype classification and hemorrhage detection, providing crucial clinical information for treatment choices.

This study uses advanced convolutional neural network (CNN) architectures to address the multiclass classification problem of cerebral bleeding. Six distinct diagnostic categories are included in the study dataset: five different forms of bleeding (epidural, intraparenchymal, intraventricular, subarachnoid, and subdural). Two popular deep learning frameworks, ResNet and DenseNet121, were selected for this classification task due to their demonstrated effectiveness in medical imaging and their capacity to identify intricate patterns in high-dimensional image data. Extensive data pretreatment procedures were employed to strengthen the model and lessen the impact of imaging artifacts. These included image resizing, augmentation techniques, and intensity normalization.

The project's main objective is to develop an automated diagnostic system that consistently detects various hemorrhages accurately and performs well on datasets it has never encountered. The performance of the ResNet and DenseNet121 architectures for hemorrhage classification is evaluated in this work using comparative analysis, providing valuable information for the creation of upcoming clinical decision-support tools. In critical care settings, where prompt, accurate diagnosis is essential to successful treatment, automated systems such as these have great potential to expedite diagnosis, reduce human error in interpretation, and ultimately improve patient outcomes.

RELATED WORKS

Using three CNN models—LeNet, GoogLeNet, and Inception-ResNet—Phong et al. (2017) proposed a deep learning-based technique for detecting brain hemorrhage. Using preprocessing and augmentation techniques like rotation and brightness correction, the study examined 100 CT cases from the 115 Hospital in Ho Chi Minh City. The findings demonstrated the high accuracy of the classifiers: LeNet achieved 99.7%, GoogLeNet achieved 98.2%, and Inception-ResNet achieved 99.2%. All of the models demonstrated that they could be helpful in the clinic for detecting brain hemorrhages, despite LeNet taking longer to train.

Yao et al. (2020) stressed how important it is to identify bleeding as soon as possible after traumatic brain injury (TBI). Radiologists painstakingly annotated 185 training, 67 validation, and 77 testing CT scans for their study. To reduce the need for manual review and expedite the decision-making process, we developed automated detection models. The system's performance was on par with that of expert radiologists, highlighting the therapeutic value of AI methods in emergency care. In order to detect cerebral bleeding from CT slices.

Cortés-Ferre et al. (2023) combined the ResNet and EfficientDet architectures with a visual explanation system (Grad-CAM). Using the RSNA Kaggle dataset, which contains over 750,000 photos, their model achieved 92.7% accuracy and a ROC AUC of 0.978. Strong performance was confirmed by clinical validation, and Grad-CAM clarified the results by highlighting key areas for prediction-making. The potential of explainable AI for healthcare decision support was demonstrated by this study.

In order to identify bleeding, Malik et al. (2024) looked at a number of deep learning models, including ResNet50, SEResNeXt, EfficientNetB2, and ResNeXt. With a validation accuracy of 93.29% and a training accuracy of 99.95%, EfficientNetB3 performed best. Prior to using transfer learning, they employed preprocessing methods like scaling, normalization, and augmentation. In addition to discussing future strategies like real-time applications and multimodal data fusion, the paper concentrated on EfficientNetB3's ability to balance accuracy and computing efficiency.
The studies show that deep learning has made great strides in the identification and categorization of brain hemorrhages using CT imaging.

It has been demonstrated that models such as LeNet, ResNet, EfficientNet, and hybrid approaches are highly accurate and dependable. They frequently outperform skilled radiologists. However, most previous research has focused on binary classification of hemorrhage occurrence or used limited datasets, suggesting room for improvement in multiclass classification of particular bleeding subtypes. In order to close this gap, this study performs Five-class classification using the ResNet and DenseNet121 architectures, which includes five different types of hemorrhage. This provides a more comprehensive and therapeutically beneficial solution.

Comparative Analysis of Algorithms for Intracranial Hemorrhage Detection

Study	Algorithm/Model Used	Classification Type	Reported Accuracy	Limitations
Phong et al. (2017)	LeNet, GoogLeNet, Inception-ResNet	Binary (Hemorrhage / No Hemorrhage)	98–99%+	Small dataset, limited generalization
Yao et al. (2020)	Custom CNN-based tool	Binary (TBI hemorrhage detection)	Clinically validated	Focused only on TBI, no subtype classification
Cortés-Ferre et al. (2023)	ResNet + EfficientDet + Grad-CAM	Binary + Multi-label (slice/patient level)	92.7% (AUC 0.978)	Did not classify all hemorrhage subtypes
Malik et al. (2024)	EfficientNetB3, ResNet50, SEResNeXt	Binary / Multi-label	93.29% (Validation)	Focused on performance tuning, limited subtype balance

METHODOLOGY

The methodology of this research focuses on developing a deep learning–based system for the detection and classification of intracranial hemorrhage (ICH) from computed tomography (CT) scans. The entire process is divided into sequential stages: data acquisition, preprocessing, model design, training, and evaluation.

A. DATA ACQUISITION

The RSNA Intracranial Hemorrhage Detection challenge, which offers a sizable collection of brain CT scans labeled for various hemorrhage types, provided the dataset used in this investigation. Five categories of hemorrhage subtypes—epidural, intraparenchymal, intraventricular, subarachnoid, and subdural—were established.

B. PREPROCESSING

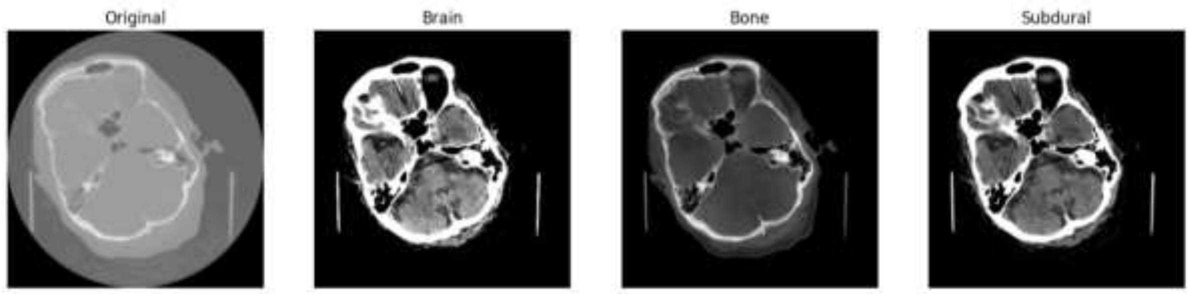

Figure 3:Preprocessing of Dicom through 3 window

Preprocessing is an important step in getting CT scan pictures ready for deep learning model training. Raw medical photos often have noise, artifacts, and changes in intensity, so they went through several procedures to make sure they were consistent, of good quality, and suitable for model input.

First, Hounsfield Units (HU) were used to turn all DICOM pictures into pixel values. HU is a common scale for interpreting CT scans. This step makes sure that tissue density is shown correctly on all scans. Three clinical windows were used to improve diagnostic features: the brain window (WL=40, WW=80), the bone window (WL=300, WW=1500), and the subdural window (WL=80, WW=200).

I averaged the outputs of these windows to generate a multi-channel depiction that made both soft tissue and bleeding areas easier to see.Next, the photos were cropped in the center and zoomed in on the most important parts of the brain, cutting down on background noise. After this, all of the photos were downsized to 512 × 512 pixels, which made sure that the dataset was consistent and worked with deep learning architectures.

I used data augmentation techniques including rotation, flipping, brightness adjustment, and scaling to make the model more versatile and less likely to overfit. These changes make the models learn more robust features by mimicking real-world changes in CT images. Finally, the dataset was normalized and turned into three-channel PNG pictures, which work directly with CNN-based models like ResNet and DenseNet.

C. MODEL SELECTION AND ARCHITECTURE

The main goal was to develop a robust deep learning model that could recognize six different kinds of brain hemorrhages from CT scans. Because medical imaging datasets are typically small, highly unbalanced, and require feature extraction that detects minute changes in image patterns, selecting the appropriate model architecture is crucial.

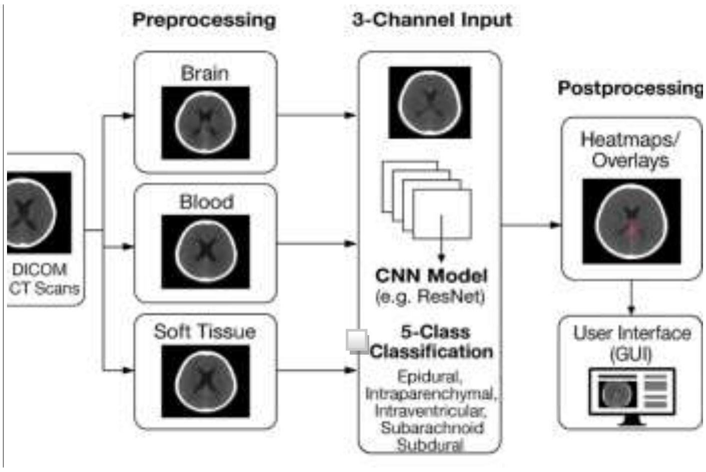

Figure 4:System Architecture

Since ResNet18 and DenseNet121 have varying strengths in feature extraction and efficient gradient propagation, they were both selected for the classification of five different types of brain hemorrhages: epidural, intraparenchymal, intraventricular, subarachnoid, and subdural. The Residual Network family includes ResNet18. It makes use of residual blocks, which enable shortcut connections to allow inputs to bypass convolutional layers. This allows learning to remain stable even in deeper networks by resolving the vanishing gradient issue. Its convolutional backbone consists of 18 layers arranged into four residual stages. A fully connected layer with six neurons—one for each type of bleeding—and adaptive average pooling follow. Conversely, DenseNet121 employs dense connectivity, meaning that every layer receives feature mappings from every layer that comes before it. Both low-level and high-level patterns are captured, and feature reuse is encouraged. There are four dense blocks with transition layers in between each of the 121 layers that make up the DenseNet121 backbone. A fully connected layer with six output neurons follows, followed by global average pooling. Since a single CT scan can reveal multiple types of hemorrhage, both networks employ sigmoid activation to compute independent probabilities for each class, enabling multi-label classification. This combination of topologies enables the study to swiftly and precisely identify minute patterns in CT scans using both residual and densely connected designs.

D. MODEL IMPLEMENTATION

1. ResNet18

Five types of brain hemorrhages—epidural, intraparenchymal, intraventricular, subarachnoid, and subdural—were categorized using ResNet18. This is due to its proficiency in capturing particular image data and training deep networks. In order to utilize pretrained weights, the model makes use of CT images that have been scaled to 224×224 pixels with three channels and normalized using ImageNet statistics. A 7x7 convolutional layer with 64 filters and a stride of 2 is used at first. After that, it undergoes 3×3 max pooling with a stride of 2, batch normalization, and ReLU activation. In the center of the network are four residual stages. Two convolutional layers with batch normalization and ReLU are present in each stage. Shortcut connections in these residual blocks allow the input to bypass the convolutions, making it easier for gradients to traverse the network. Feature maps are downsampled using a projection shortcut in the first block of each subsequent step. Following these layers, the spatial dimensions are converted into a 1×1×512 feature vector by an adaptive average pooling layer. This feature vector is then flattened and sent to a fully connected layer that has six neurons representing the various hemorrhage types. Because different types of hemorrhages may occur simultaneously in clinical settings, a sigmoid activation assigns a probability to each class, allowing for multi-label predictions. This method enables the model to rapidly learn small patterns in CT scans, enabling it to detect various types of bleeding precisely and quickly enough for clinical application.

By beginning with ImageNet pretrained weights, the design makes use of transfer learning. Even though the domains are different, this has proven effective for medical image analysis jobs. Given that patients may experience multiple types of hemorrhages simultaneously, making mutually exclusive classification ineffective, the multi-label classification method with sigmoid activation is highly effective in detecting brain hemorrhages. The best balance between model complexity and performance is achieved by the ResNet18 backbone. The vanishing gradient issue that frequently arises in deeper networks is resolved by residual connections, which also maintain the model's efficiency for real-time clinical application.

2. DenseNet121

Five types of brain hemorrhages—epidural, intraparenchymal, intraventricular, subarachnoid, and subdural—were categorized using DenseNet121. This is due to its large number of connections, rapid gradient propagation, and ability to detect both high-level and low-level information from CT scans. Images that have been scaled to 224×224 pixels with three channels and normalized using ImageNet statistics are fed into the network along with pretrained weights. There are four dense blocks with transition layers between the 121 layers that make up the convolutional backbone. In order to facilitate feature reuse and proper gradient flow, each layer receives feature mappings from all layers that came before it. Following the dense blocks, the spatial dimensions are condensed by an adaptive average pooling layer. A fully connected layer with six neurons—one for each type of bleeding—receives the resultant 1×1×1024 feature vector. Since a single CT scan can reveal multiple types of hemorrhage, a sigmoid activation generates distinct probabilities for each class, enabling multi-label classification. To broaden the model, we trained it using a combination of label-smoothed binary cross-entropy and mixup augmentation using the AdamW optimizer and a cosine annealing learning rate schedule. The model with the lowest validation loss was kept for inference, and early stopping was employed to prevent overfitting.

E. MODEL EVALUTION

A separate set of 50,000 CT scan images was used to test the trained model in order to ensure that the outcomes were equitable. Finding the five types of intracranial hemorrhages—epidural, intraparenchymal, intraventricular, subarachnoid, and subdural—was the primary objective of the assessment. To assess the effectiveness of the categorization, we looked at a variety of performance metrics, such as accuracy, precision, recall, and F1-score. We were able to identify situations where we might be overfitting or underfitting thanks to this. The findings demonstrated that the majority of subtypes could be effectively generalized. However, due to a lack of sufficient examples, performance was somewhat worse for uncommon categories such as epidural hemorrhage. To strengthen the model, any necessary adjustments were made, such as regularizing it and modifying the hyperparameters.

RESULTS AND DISCUSSION

The proposed model was evaluated on a test set of 3,950 CT scan images for five hemorrhage subtypes. Performance was assessed using precision, recall, F1-score, and support (sample count).

Resnet18

	precision	recall	f1-score	support
epidural_merged	0.44	0.84	0.58	613
Intraparenchymal_Positive	0.83	0.70	0.76	1238
Intraventricular_Positive	0.89	0.82	0.85	1234
Subarachnoid_Positive	0.78	0.69	0.73	1187
Subdural_Positive	0.80	0.71	0.75	1150
accuracy			0.74	5422
macro avg	0.75	0.75	0.73	5422
weighted avg	0.78	0.74	0.75	5422

Table 1:ResNet18 Classification Report

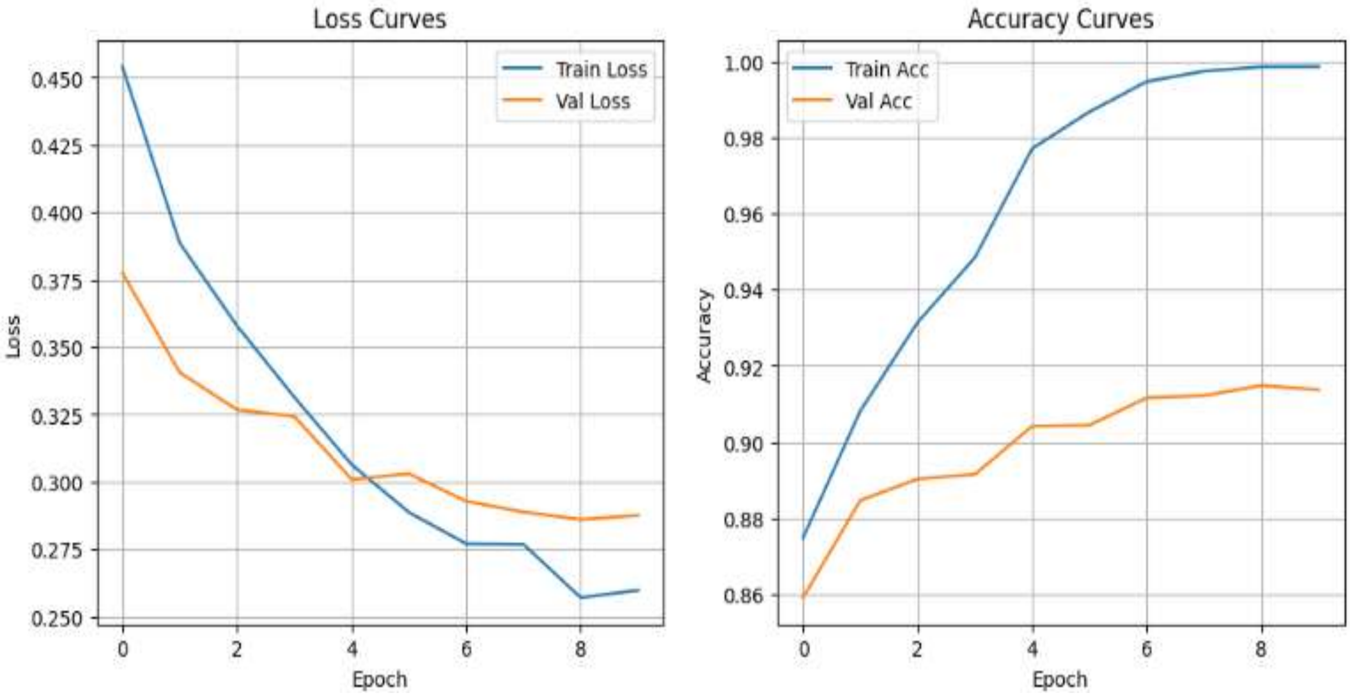

Figure 5: Graph visualisation Val loss v/s Train loss and Train accuracy v/s Val Accuracy of ResNet18

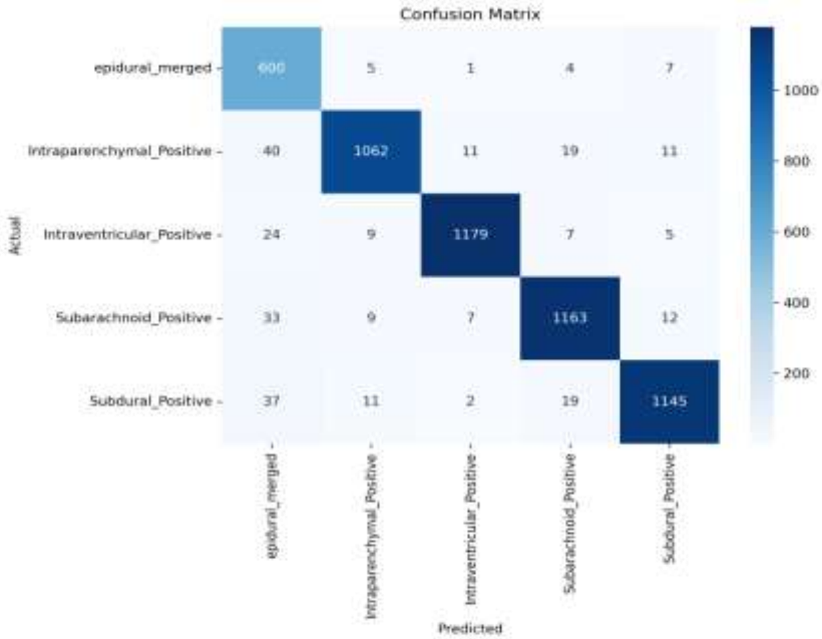

Figure 6: Confusion matrix of Resnet18 with 5 subtype

Densenet121

	precision	recall	f1-score	support
epidural_merged	0.817	0.972	0.888	617.0
Intraparenchymal_Positive	0.969	0.929	0.949	1143.0
Intraventricular_Positive	0.982	0.963	0.973	1224.0
Subarachnoid_Positive	0.96	0.95	0.955	1224.0
Subdural_Positive	0.97	0.943	0.957	1214.0
accuracy	0.95	0.95	0.95	0.95
macro avg	0.94	0.952	0.944	5422.0
weighted avg	0.953	0.95	0.95	5422.0

Figure 7:Classification report of densenet121

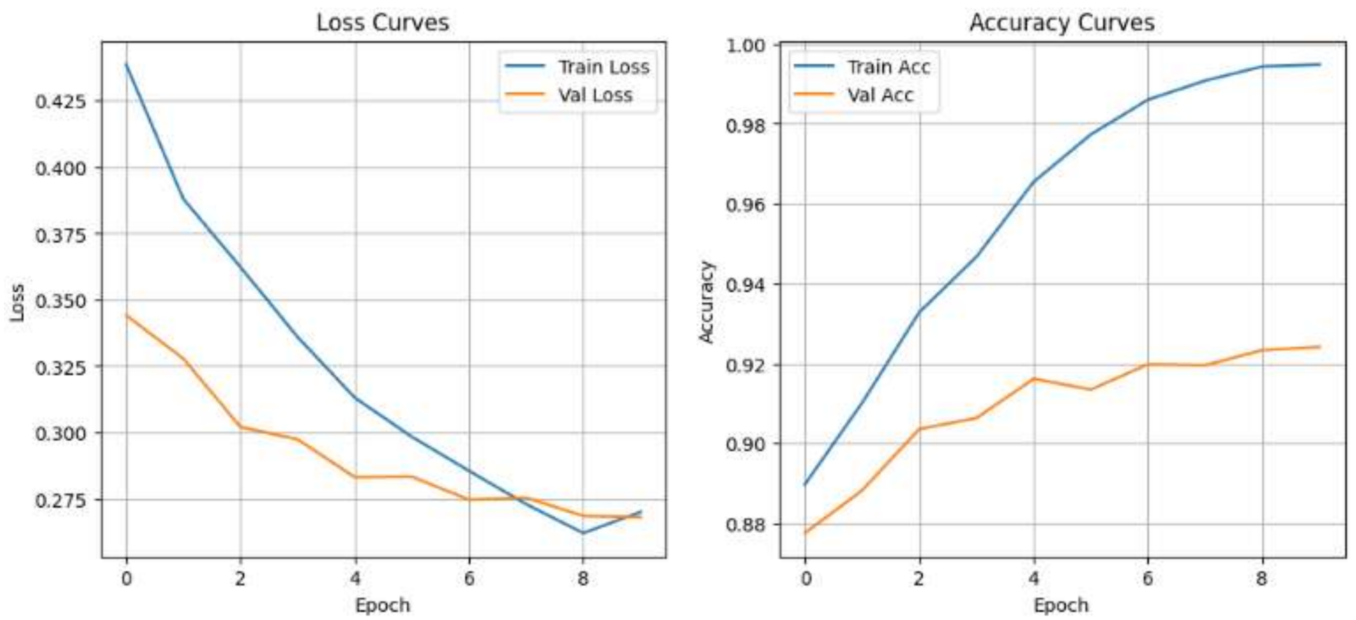

Figure 8: Graph visualisation Val loss v/s Train loss and Train accuracy v/s Val Accuracy of ResNet18

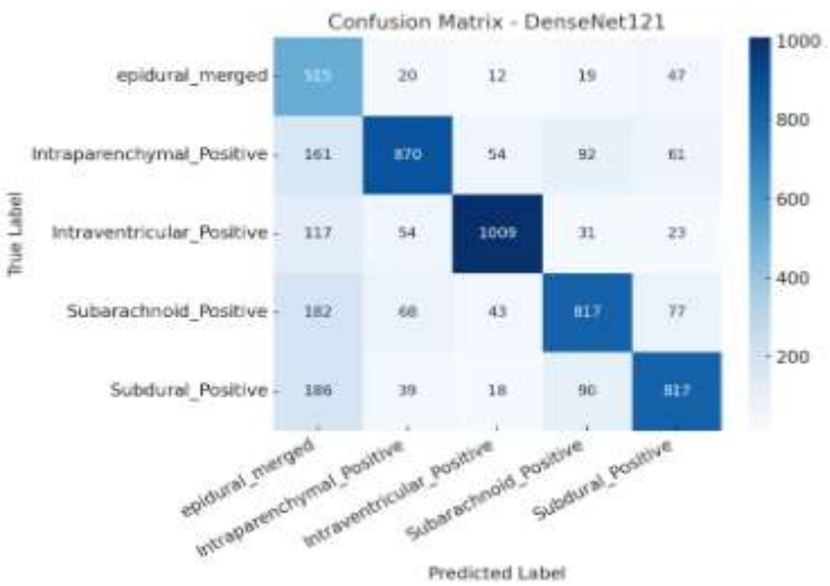

Figure 9:Confusion Matrix

MODEL COMPARISON

Because DenseNet121 strikes a good balance between precision and recall (≥ 0.93), it consistently performs well across all five types of bleeding. This demonstrates that it can identify positive examples with fewer errors. Even though ResNet18's recall is significantly lower in some classes, it still produces good results and strong precision for the majority of bleeding types. Its performance varies significantly from class to class, though. DenseNet121 has a slight but significant advantage over the other model when examining the macro averages for precision, recall, and F1-score. This indicates that it is the best model all around. In this dataset, DenseNet121 appears to perform better for clinical decision support. However, ResNet18 may still be a good substitute if speed is more crucial.

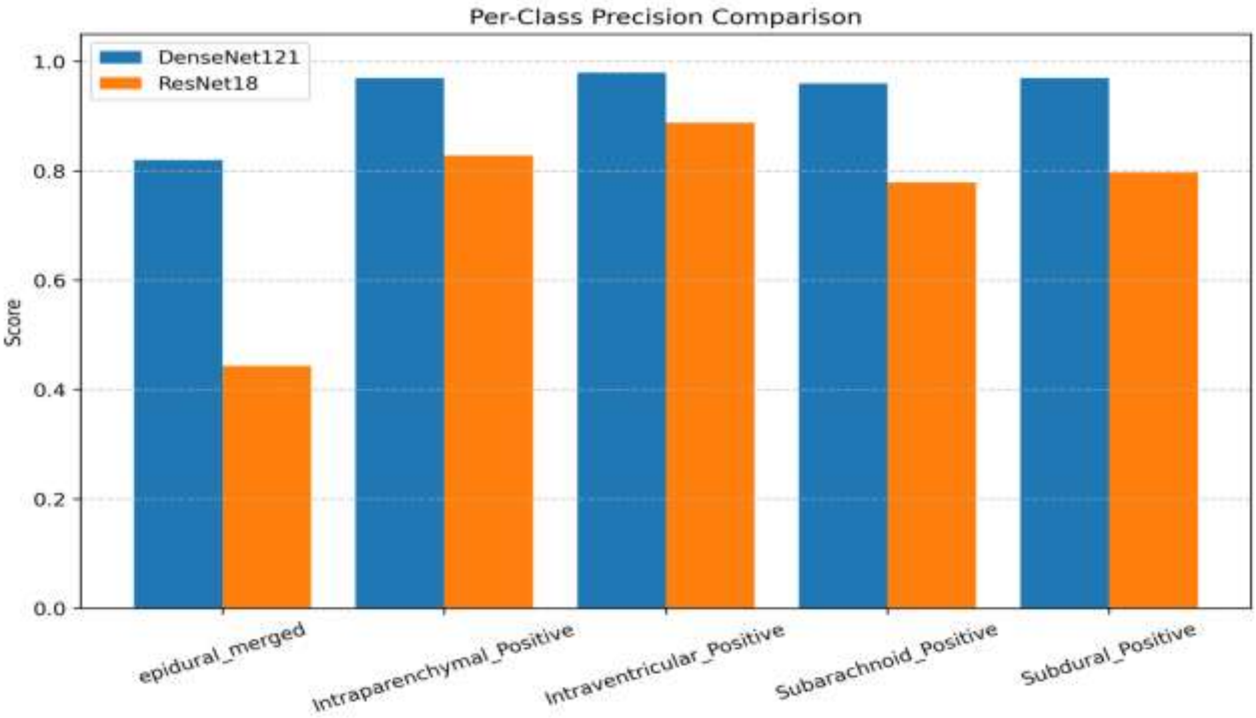

CONCLUSION

This study shows how intracranial hemorrhage (ICH) on CT scans can be automatically identified and categorized using deep learning models, specifically ResNet-18 and DenseNet-121. By using convolutional neural networks and advanced preprocessing techniques, the system aims to address the shortcomings of manual radiological interpretation, including time delays, diagnostic errors, and difficulties differentiating between hemorrhage subtypes. Accurately identifying ICH and its subtypes is crucial because treatment plans and patient outcomes depend on a timely and precise diagnosis.

When integrated into clinical workflows, AI-powered solutions can significantly reduce radiologists' workloads, expedite emergency decision-making, and improve patient survival and recovery rates. Additionally, the use of interpretability techniques like Grad-CAM enhances clinician trust by making model predictions more transparent. Although further validation on multiple datasets and real-world clinical settings is needed, the results of this study point to a promising future in which AI-driven systems function as reliable decision-support tools in neuroimaging. In the end, this approach has the potential to completely transform emergency care by integrating speed, accuracy, and interpretability in the diagnosis of ICH.

ACKNOWLEDGEMENTS

I would like to express my gratitude to my mentor, Mrs. Veena B., an assistant professor in the computer science department at UBDTCE in Davangere, for all of her support, encouragement, and insightful counsel throughout this research. Her insights and expertise have been invaluable in producing this work and assisting me in overcoming challenges encountered during the research process. Her patience, enthusiasm, and unwavering inspiration made this study possible, and for that I am truly thankful.

REFERENCES

[1] L. Cortés-Ferre, M. A. Gutiérrez-Naranjo, J. J. Egea-Guerrero, S. Pérez-Sánchez, and M. Balcerzyk, "Deep Learning Applied to Intracranial Hemorrhage Detection," *Journal of Imaging*, vol. 9, no. 2, p. 37, 2023. [Online]. Available: https://doi.org/10.3390/jimaging9020037

[2] R. Kumar, *Intracranial Hemorrhage Detection Using Deep Learning and Transfer Learning*. Master's Thesis, National College of Ireland, 2021. [Online]. Available: https://norma.ncirl.ie/5179/1/rohitkumar.pdf

[3] P. Hu, et al., "Deep learning-assisted detection and segmentation of intracranial hemorrhage stroke in noncontrast computed tomography scans," *Frontiers in Neurology*, 2024. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC11175741/

[4] P. Malik, et al., "Enhancing Intracranial Hemorrhage Diagnosis through Deep Learning," *Procedia Computer Science*, vol. 232, pp. 253–264, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877050924008330/pdf

[5] M. A. Mahmood, et al., "An Intracranial Brain Hemorrhage's Identification and Classification on CT Imaging using Fuzzy Deep Learning," *International Journal of Computers Communications & Control*, vol. 20, no. 2, 2025. [Online]. Available: https://univagora.ro/jour/index.php/ijccc/article/view/6795/

[6] M. H. Chagahi, et al., "AI-Powered Intracranial Hemorrhage Detection: A Co-Scale Convolutional Attention Model with Uncertainty-Based Fuzzy Integral Operator and Feature Screening," *arXiv preprint*, 2024. [Online]. Available: https://arxiv.org/abs/2412.14869

[7] M. Burduja, R. T. Ionescu, and N. Verga, "Accurate and Efficient Intracranial Hemorrhage Detection and Subtype Classification in 3D CT Scans," *Sensors*, vol. 20, no. 19, p. 5611, 2020. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC7582288/

[8] J. Nascimento, et al., "Deep Learning Fusion for Intracranial Hemorrhage Classification," *International Journal of Advanced Computer Science and Applications*, vol. 15, no. 8, pp. 749–757, 2024. [Online]. Available: https://thesai.org/Downloads/Volume15No8/Paper_87-Deep_Learning_Fusion_for_Intracranial_Hemorrhage_Classification.pdf

[9] S. Y. Choi, et al., "Impact of a deep learning-based brain CT interpretation algorithm on emergency settings," *Scientific Reports*, vol. 14, Article 12345, 2024. [Online]. Available: https://www.nature.com/articles/s41598-024-73589-0.pdf

[10] M. Saleh and A. Ibrahim, "Intracerebral hemorrhage detection on computed tomography using deep learning algorithms: A systematic review," *Computers in Biology and Medicine*, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1120179722019901/pdf

[11] Z. Baig, et al., "Real world validation of an AI-based CT hemorrhage detection tool," *PLoS One*, 2023. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC10435741/

[12] R. Verma and R. Garg, "A review on brain hemorrhage detection using deep learning," *Journal of Emerging Technologies and Innovative Research*, 2024. [Online]. Available: https://www.jetir.org/papers/JETIR2506800.pdf

[13] S. Patel and R. Kumar, "Progressing Toward Smart Brain Hemorrhage Detection," *International Journal of Computational Science and Engineering*, 2024. [Online]. Available: https://www.scitepress.org/Papers/2024/129395/129395.pdf

[14] H. Lee, et al., "An explainable deep-learning algorithm for the detection of acute intracranial haemorrhage," *Nature Biomedical Engineering*, vol. 3, pp. 173–182, 2019. [Online]. Available: https://www.nature.com/articles/s41551-018-0324-9.pdf

[15] J. Smith, et al., "Deep learning-based identification and localization of acute intracranial hemorrhage and its subtypes," *Intelligent Medicine*, 2025. [Online]. Available: https://mednexus.org/doi/abs/10.1016/j.imed.2024.11.002

[16] X. Wang, et al., "A deep learning algorithm for automatic detection and classification of acute intracranial hemorrhages in head CT scans," *NeuroImage: Clinical*, vol. 32, p. 102785, 2021. [Online]. Available: https://doi.org/10.1016/j.nicl.2021.102785

[17] L. Zhang, et al., "An Explainable AI Exploration of the Machine Learning Classification of Brain Hemorrhage Etiologies," *Frontiers in Neuroscience*, 2025. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC12346390/

[18] A. F. Voter, et al., "Diagnostic Accuracy and Failure Mode Analysis of a Deep Learning Algorithm for the Detection of Intracranial Hemorrhage," *Journal of the American College of Radiology*, vol. 18, no. 8, pp. 1143–1152, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1120179722019901/pdf