Received Date:22nd December, 2024
Revision Date: 1st January, 2025
Accepted Date:9th January, 2025

Detection of Intracranial Hemorrhage Using Deep Learning

Sahayog Timalsina[1*], Rabindra Joshi[2], Siddhant Bhatta[3], Samir Pokharel[4], Shyam Dahal[5]
[1]*Dept. of Computer Engineering, Kathmandu Engineering College, E-mail: sahayogtimalsina07@gmail.com*
[2]*Dept. of Computer Engineering, Kathmandu Engineering College, E-mail: joshirabindra2058@gmail.com*
[3]*Dept. of Computer Engineering, Kathmandu Engineering College, E-mail: bhattasiddhant950@gmail.com*
[4]*Dept. of Computer Engineering, Kathmandu Engineering College, E-mail: pokharelsamir246@gmail.com*
[5]*Assoc. Professor, Dept. of Computer Engineering, Kathmandu Engineering College, E-mail: shyam.dahal@kecktm.edu.np*

Abstract — Intracranial hemorrhage has always been a crucial medical condition where bleeding within the cranium region occurs, leading to severe neurological damage and sudden demise of a person. A patient's likelihood of survival in the treatment of Intracranial Hemorrhage is dependent on rapid diagnosis based on the radiologist's assessment of Computed Tomography (CT) scans. As a result, the requirement for precise and prompt identification of Intracranial Hemorrhage (ICH) is of utmost importance. Deep learning models can be used to assist this process by accelerating the time it takes to identify them. We built a deep learning model which will accelerate the time required to identify intracranial hemorrhages such that it facilitates the classification and segmentation of Intracranial Hemorrhage. We have constructed an EfficientNetB4 model using a Convolutional Neural Network (CNN) architecture which was used for the classification of images. This model comprises various layers categorized into different types, including convolution layers, normalization layers, fully connected layers, and activation layers. And we used Grad-Cam model for the image segmentation process which generated heat map for the image that contained the hemorrhage. We have used approximately 700,000 DICOM files collected from four international universities by the Radiological Society of North America (RSNA). We achieved an accuracy of 97 percent with learning rate of 0. 000125, batch size of 32 and 15 epochs of model training.

Keywords — Intracranial Hemorrhage, Computed Tomography, Deep Learning

Introduction

Intracranial hemorrhage (ICH), or brain hemorrhage, occurs when blood vessels within the skull or brain rupture, leading to bleeding. Hypertension, vascular abnormalities such as aneurysms or arteriovenous malformations (AVMs), trauma, blood thinners, and spontaneous factors are common causes. Existing literature reviews explore the use of deep learning for detecting ICH from CT scans, highlighting methodologies, challenges, and advancements in this area. However, there is currently no published system utilizing deep learning for ICH detection. In Nepal, ICH is typically detected using CT scans or MRI machines, where radiologists examine images to identify and assess hemorrhages, informing treatment plans. However, limitations exist, including a shortage of trained radiologists and the time-consuming nature of manual classification, which may impact patient outcomes. To address these challenges, we propose a system designed to classify hemorrhages based on CT scans. Our system offers several key benefits. Firstly, it enables early detection of ICH by leveraging a deep learning model capable of identifying subtle signs of hemorrhage that may be difficult for radiologists to detect initially. Early detection facilitates prompt medical intervention, potentially improving patient outcomes. Additionally, our system automates the detection and segmentation process, reducing the need for manual intervention by radiologists. By automating these tasks, our system saves time and effort for medical professionals, allowing them to focus on other critical aspects of patient care. In summary, our proposed system aims to address the limitations of existing methods for detecting intracranial hemorrhage by leveraging deep learning technology. Through early detection and automation of the detection process, we believe our system has the potential to improve patient outcomes and streamline medical workflows in the diagnosis and treatment of ICH.

Related papers

1. A study from Vingroup Big Data Institute, Hanoi, Vietnam, introduces a new method for intracranial hemorrhage detection on CT scans. They combine CNN with LSTM for precise prediction, training the entire architecture end-to-end. RGB-like images,

** Corresponding Author*

created by stacking three viewing windows of a single slice, serve as input. Validation was performed on the RSNA Intracranial Hemorrhage Detection challenge and the CQ500 dataset. Utilizing advanced pre-trained models like ResNet and SE-ResNeXt, known for their effectiveness on natural images, they extract 2048 features per input slice. Comparison of their models reveals SE-ResNeXt's superior accuracy, determined by weighted log loss [1].

2. An IEEE study explores intracranial hemorrhage detection using a deep CNN to learn features and classify hemorrhages in CT scans. The dataset comprises 134 CT cases with approximately 1.1 billion voxels, divided for training, validation, and testing. Performance on the test set shows 81% sensitivity per lesion and 98% specificity per case. An annotation tool in MATLAB aids in marking hemorrhage boundaries. However, despite high per-voxel specificity (0.935), false alarms occur in every test case [2].

3. In a separate study, researchers propose techniques based on ResNet CNN and EfficientDet to enhance deep learning models for digital image detection. They incorporate the Gradient-weighted Class Activation Mapping (Grad-CAM) method to offer visual explanations through gradient-based localization. Utilizing both Kaggle dataset and external data for training, validation, and testing, their aim is to introduce an integrated deep learning model for intracranial hemorrhage detection in brain CT scans. This model, along with a visual explanation system, aims to aid human experts in decision-making by detecting bleeding in DICOM 2D image files [3].

4. The study evaluated the performance of different models on original and preprocessed brain CT images for classifying intracranial hemorrhage (ICH) and normal cases. The modified AlexNet- SVM classifier showed superior performance, achieving specificity and recall rates of 0.9986 for original images and 0.9953 and 0.9967, respectively, for preprocessed images. All models achieved an AUC > 0.9996 for binary classification. Overall, the research presents a simple and effective framework for classifying hemorrhage and non-hemorrhage images, serving as a valuable screening tool for assisting radiological trainees in accurate ICH detection [4].

Methodology

A. System Architecture

1. *EfficientNet:*

EfficientNet, introduced in 2019 by Mingxing Tan and Quoc V. Le, represents a breakthrough in convolutional neural network (CNN) architecture by offering state-of-the-art accuracy with fewer parameters and computational resources compared to counterparts like ResNet or Inception. Its core innovation lies in a compound scaling strategy, which adjusts the network along three dimensions - depth, width, and resolution - to strike a balance between accuracy and computational efficiency.

Depth scaling involves increasing the number of layers in the network to capture more intricate features, while width scaling boosts the number of channels in each layer to diversify feature extraction. Resolution scaling adjusts the input image size to balance detail and computational demand.

EfficientNet achieves this by employing a compound coefficient, φ, which uniformly scales network width, depth, and resolution in a systematic manner. This is mathematically represented as:

Depth $d = \alpha\varphi$, Width $w = \beta\varphi$, Resolution $r = \gamma\varphi$, (1) such that $\alpha.\ \beta 2.\ \gamma 2 \_ 2$ $\alpha >= 1$, $\beta >= 1$, $\gamma >= 1$ where α, β, and γ are determined through empirical validation via a grid search on a validation set. These coefficients ensure that the overall model size remains efficient while maximizing accuracy. The values of α, β, and γ are determined through a grid search algorithm, ensuring a constant ratio of scaling across width, depth, and resolution. By varying the compound coefficient φ, EfficientNet can generate models ranging from B1 to B7, each optimized for different computational resources. This approach represents a significant advancement in CNN architecture design, allowing for unparalleled efficiency without sacrificing performance.

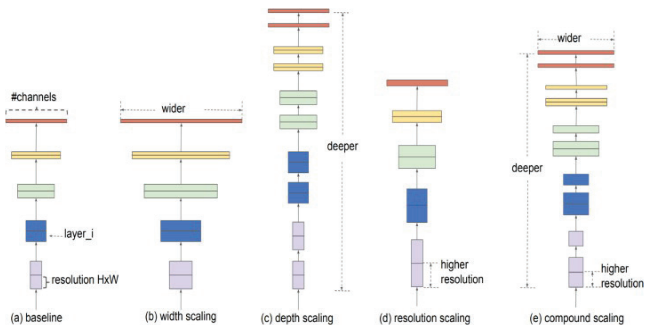

Figure 1: Model Scaling [5]

2. Transfer Learning

Transfer learning means borrowing knowledge from a pre-trained model and applying it to a new task instead of starting from scratch. This saves time and resources. Pre- trained models, like those trained on ImageNet for recognizing objects, already understand basic patterns. You can use them directly for new tasks or tweak them a bit. There are two main ways to do this: feature extraction and fine-tuning. Feature extraction keeps the basic pattern recognition intact and only adjusts the final part of the model. Fine-tuning adjusts both the basic patterns and the final part. We chose fine-tuning because our dataset involves things not in ImageNet. We tweaked the model by changing some layers to recognize our specific features, like hemorrhages, and trained it to make accurate predictions.

3. Grad-CAM:

Grad-CAM, short for Gradient Weighted Class Activation Map, identifies the image areas the model focuses on during classification. It combines gradient information with class activation maps, making it adaptable to various networks. The process involves finding the last convolutional layer and monitoring gradient information. It generates a heat map highlighting pixels relevant to classification. Grad- CAM is straightforward, compatible with any CNN architecture, and offers insights into model focus. In our project, it aids in segmenting hemorrhage regions in images. If no hemorrhage is detected, the original image is retained.

B. Designs

i. Dataset

The RSNA Intracranial Hemorrhage Detection dataset is a comprehensive collection of CT scans used to identify acute intracranial hemorrhage and its subtypes. This dataset was sourced from a Kaggle competition organized by the Radiological Society of North America (RSNA) and comprises 752,803 images, all provided in DICOM format. DICOM images include associated metadata such as Patient ID, StudyInstanceUID, and SeriesInstanceUID, which are essential for organizing and analyzing the data. The dataset is labeled with six categories of hemorrhage: epidural, intraparenchymal, intraventricular, subarachnoid, subdural, and a general category labeled "any," which is true if any type of hemorrhage is present. This dataset serves as the foundation for training and evaluating our intracranial hemorrhage detection model.

ii. Preprocessing

Preprocessing is a critical step in preparing raw data for analysis, modeling, or machine learning. It involves cleaning, organizing, and enhancing data quality by addressing missing values, reducing noise, and standardizing data representation. In our project, several preprocessing techniques were employed to optimize the dataset for model training:

1. *DICOM Windowing:* This technique adjusts the contrast and brightness of medical images to improve the visualization of anatomical structures. By selecting specific ranges of pixel values and mapping them to a display window, DICOM windowing enhances the clarity of CT scans, making it easier to identify hemorrhages.

2. *Image Augmentation:* To increase the diversity of the training data, we applied transformations such as random flips (horizontal with a 25% probability and vertical with a 10% probability). This technique helps the model generalize better by exposing it to varied representations of the same data.

   These preprocessing steps ensure that the input data is in a suitable format for training the model, improving both data quality and model performance.

iii. Convolutional Neural Networks (CNN)

Convolutional Neural Networks (CNNs) are a specialized type of deep learning architecture designed for image analysis and recognition tasks. They are particularly effective in computer vision applications such as image classification, object detection, and segmentation. The key components of a CNN include:

1. *Input Layer:* This layer receives the input data, typically in the form of a multidimensional array representing pixel intensity or color values.

2. *Convolutional Layer:* This layer extracts features from the input data using filters (kernels) that identify patterns such as edges, textures, and shapes. The output is a set of feature maps that highlight important characteristics of the input.

3. *Activation Function:* Non-linearity is introduced using activation functions like ReLU (Rectified Linear Unit), which sets negative values to zero and keeps positive values unchanged.

4. *Pooling Layer:* This layer reduces the spatial dimensions of the feature maps while retaining essential information. Common pooling operations include max pooling and average pooling.
5. *Fully Connected Layer:* This layer connects every neuron in one layer to every neuron in the next, enabling the network to learn high-level features from the extracted data.
6. *Output Layer:* The final layer produces the network's predictions or classifications based on the task, such as identifying hemorrhage types in CT scans.

CNNs are highly effective for medical image analysis due to their ability to learn hierarchical representations of data, making them ideal for tasks like intracranial hemorrhage detection.

iv. Verification and Validation

For intracranial hemorrhage detection, we employed EfficientNetB4, a variant of the EfficientNet architecture known for its balance between model size and performance. EfficientNetB4 uses a compound scaling method to uniformly scale network width, depth, and resolution, achieving state-of-the-art performance with fewer parameters compared to traditional CNNs.

1. *Model Training and Fine-Tuning:* We utilized transfer learning to fine-tune the pre-trained EfficientNetB4 model over 15 epochs with a batch size of 32. This approach allowed the model to adapt its learned features to our specific task, improving convergence and performance.
2. *Optimization and Evaluation:* The model was trained with a learning rate of 0.000125 and a batch size of 32, carefully selected through experimentation to balance computational efficiency and model accuracy.
3. *Results:* The model achieved a validation loss of 0.025 and a validation accuracy of 97%, demonstrating its robustness in accurately detecting intracranial hemorrhages. These results highlight the model's potential for real-world medical applications, aiding healthcare professionals in timely and accurate diagnosis.

Result

This section demonstrates results obtained while creating our web application for analyzing Computed Tomography (CT) scan files using Efficient Net and Grad- CAM segmentation models. The graphs below illustrate the final results we achieved:

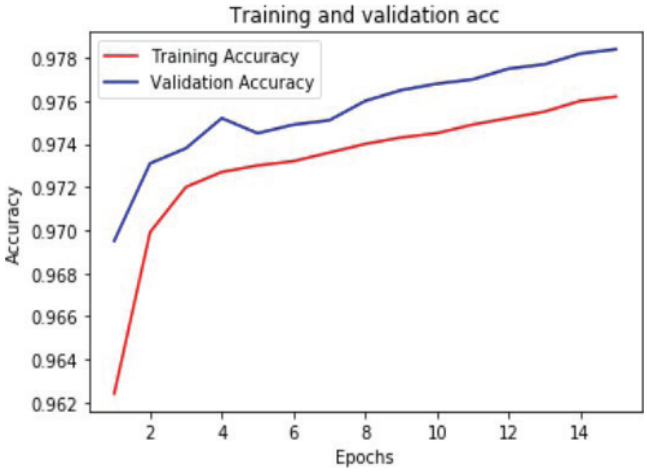

Figure 2: Training and validation accuracy

i. Training and Validation Accuracy:

The graph illustrates the training and validation accuracy over multiple epochs. The training accuracy (red line) starts lower but increases steadily, while the validation accuracy (blue line) also improves and remains slightly higher than the training accuracy throughout the epochs. This suggests that the model is learning effectively and generalizing well to unseen data without significant overfitting

ii. Training and Validation Loss:

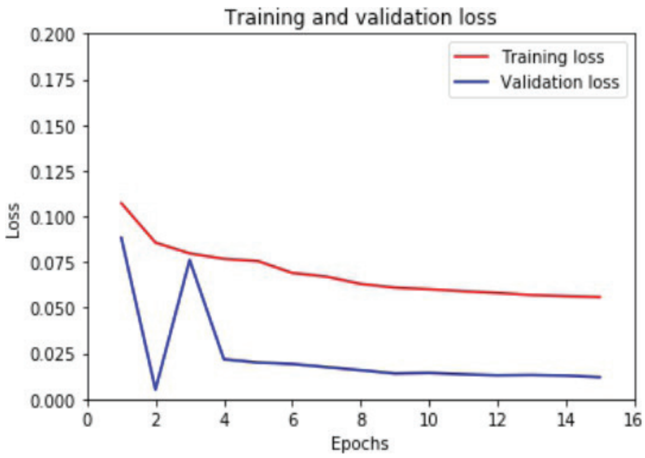

Figure 3: Training and validation loss

The graph shows the training and validation loss over the epochs. The training loss (red line) decreases gradually, indicating that the model is learning and minimizing errors. The validation loss (blue line) fluctuates initially but then stabilizes, suggesting that the model's performance on unseen data is improving while maintaining a good balance

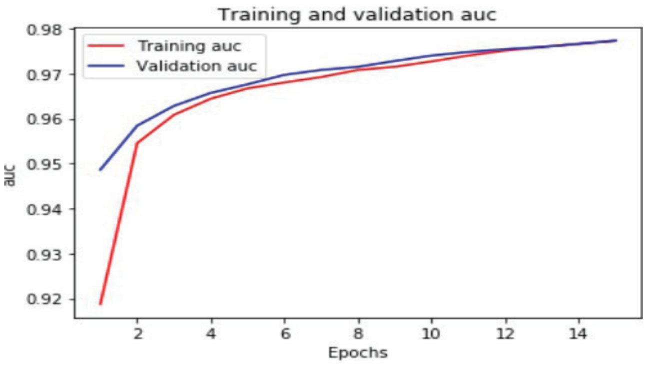

Figure 4: Training and validation AUC

iii. *Training and Validation AUC*

The graph represents the training and validation AUC (Area Under the Curve), which measures the model's ability to distinguish between classes. Both training AUC (red line) and validation AUC (blue line) increase over time, with validation AUC slightly outperforming training AUC. This indicates that the model is effectively learning and improving its classification performance over successive epochs.

The confusion matrix, bar charts undertaking just random 1506 images are shown below:

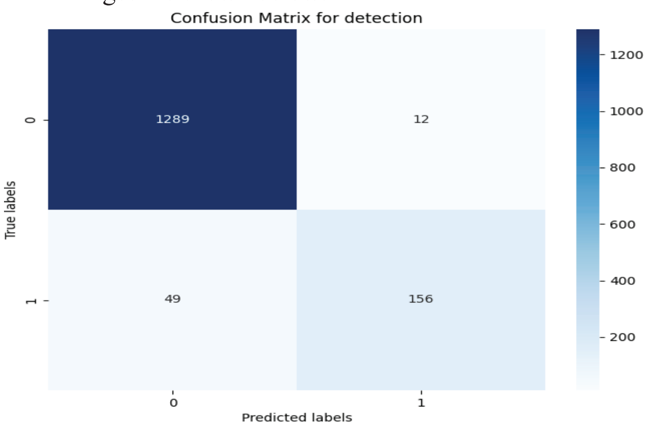

Figure 5: Confusion Matrix

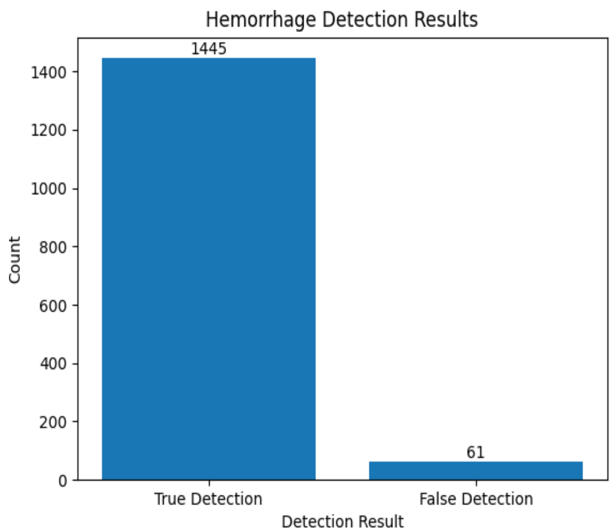

Figure 6: Bar Chart of hemorrhage detection

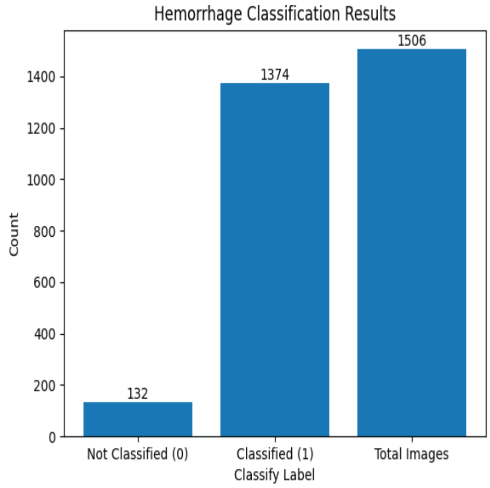

Figure 7: Bar Chart of hemorrhage classification

Conclusion

The model we developed was capable of receiving CT scan images as input from users, processing and adapting them to the necessary format for our model, and determining the presence of hemorrhages. In cases where hemorrhages are identified, the program is able to specify their type and location within the image. This project was successfully completed through the training of the EfficientNetB4 architecture on a dataset comprising approximately 700,000 DICOM files sourced from four international universities under the supervision of the Radiological Society of North America (RSNA). The resulting model achieves an accuracy rate of 97%.

Acknowledgment

We extend our heartfelt appreciation to the Department of Computer Engineering at Kathmandu Engineering College for granting us the opportunity to undertake this project. This endeavor has significantly enhanced our programming skills, benefiting our future endeavors. We are grateful to Er. Shyam Dahal, our project supervisor, and Er. Sudeep Shakya, the Head of the Computer Department, for their invaluable guidance, constructive criticism, and friendly advice, which have helped us hone our skills. Special thanks to Er. Sharad Chandra Joshi, the year coordinator and major project coordinator, for his ongoing guidance, encouragement, and valuable suggestions throughout the project duration. Additionally, we extend our sincere thanks to Dr. Sharad Gajuryal, MD Hospital Administration (MDHA), and the team at Annapurna Neurological Institute and Allied Sciences for their support during our field visit. We also express our gratitude to all our colleagues and faculty members for their direct and indirect support, cooperation, and encouragement throughout the project journey. Finally, we extend our thanks to KEC for providing us with an exceptional platform to showcase our knowledge and skills.

References

[1] A. Sage and P. Badura, "Intracranial hemorrhage detection in head CT using Double-Branch convolutional neural network, support vector machine, and random forest," Applied Sciences, vol. 10, no. 21, p. 7577, Oct. 2020, doi: 10.3390/app10217577.

[2] M. Rajagopal, S. Buradagunta, M. Almeshari, Y. Alzamil, R. Ramalingam, and V. Ravi, "An efficient framework to detect intracranial hemorrhage using hybrid deep neural networks," Brain Sciences, vol. 13, no. 3, p. 400, Feb. 2023, doi: 10.3390/brainsci13030400.

[3] 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, Feb. 2023, doi: 10.3390/jimaging9020037.

[4] P. Kumaravel, S. Mohan, J. Arivudaiyanambi, N. Shajil, and H. N. Venkatakrishnan, "A simplified framework for the detection of intracranial hemorrhage in CT brain images using deep learning," http://www.eurekaselect.com. https://www.eurekaselect.com/article/114325

[5] "Papers with Code - EfficientNet Explained." https:// paperswithcode.com /method /efficientnet