



ALZHEIMER DETECTION USING MRI IMAGES

¹ Rusheel Sable, ² Divyesh Parmar, ³ Vinay Admane, ⁴ Nishant Kasliwal ⁵ Anita Gunjal

¹Student – B.Tech. CSE, ²Student – B.Tech. CSE, ³Student – B.Tech. CSE, ⁴ Student – B.Tech. CSE ⁵
Mentor

¹Department of Computer Engineering and Technology, ² Department of Computer Engineering and
Technology,

³ Department of Computer Engineering and Technology, ⁴ Department of Computer Engineering and
Technology

⁵ Department of Computer Engineering and Technology

¹ Mit-World Peace University, Pune, India, ² Mit-World Peace University, Pune, India,

³ Mit-World Peace University, Pune, India, ⁴ Mit-World Peace University, Pune, India ⁵ Mit-World Peace
University, Pune, India

Abstract: Alzheimer's Disease (AD) is a neurodegenerative disorder that impairs memory, cognition, and daily functioning, with early diagnosis being essential for managing its progression. This research explores a deep learning-based approach to detect Alzheimer's stages from MRI scans, using several neural network architectures: Artificial Neural Networks (ANNs), Convolutional Neural Networks (CNNs), and a transfer learning model based on InceptionV3.

To address the challenge of class imbalance inherent in medical datasets, this study employs Synthetic Minority Oversampling Technique (SMOTE) and extensive data augmentation. The models classify MRI images into four categories: Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented. Experimental results demonstrate that the InceptionV3 model achieves the best performance, with significant improvements in classification metrics due to data augmentation and synthetic data generation.

This study highlights the potential of deep learning frameworks in early Alzheimer's detection, providing insights into the application of advanced machine learning techniques in medical imaging

Index Terms - Alzheimer's Disease, Artificial Neural Networks, Convolutional Neural Networks, transfer learning, InceptionV3 model, class imbalance, Synthetic Minority Oversampling Technique, SMOTE, data augmentation, CNN architecture, ANN, AUC (Area Under the Curve), softmax activation, image normalization, dropout layers, training epochs, batch

I. INTRODUCTION

Alzheimer's Disease (AD) is a progressive and irreversible neurodegenerative disorder that predominantly affects the elderly. It is the most common cause of dementia, accounting for approximately 60-70% of all dementia cases worldwide. Characterized by memory loss, cognitive decline, and changes in behavior, Alzheimer's poses significant challenges to patients, families, and healthcare systems. The widespread occurrence of Alzheimer's highlights the critical importance of early detection and timely intervention to effectively manage symptoms and slow the progression of the disease.

Early diagnosis of Alzheimer's is crucial as it allows patients to access timely treatment and participate in clinical trials that can potentially delay the onset of severe symptoms. Among the various diagnostic tools available, Magnetic Resonance Imaging (MRI) plays a pivotal role. MRI provides high-resolution images of brain structures, enabling clinicians to identify early signs of Alzheimer's, such as hippocampal atrophy and

changes in cortical thickness. However, traditional diagnostic approaches relying on manual image analysis are time-consuming, prone to subjectivity, and require specialized expertise.

The rapid advancement of deep learning techniques offers a promising solution to these limitations. Deep learning models, particularly Convolutional Neural Networks (CNNs), have shown remarkable success in automating medical image analysis and classification tasks. These models can learn complex patterns from raw image data, making them ideal for diagnosing Alzheimer's based on subtle structural changes in MRI scans. Transfer learning models, such as InceptionV3, further enhance performance by leveraging knowledge from pre-trained networks on large datasets.

Despite the potential of deep learning, one of the primary challenges in medical imaging is class imbalance. In Alzheimer's datasets, early-stage cases are often underrepresented, leading to biased models that perform poorly in detecting less common but critical stages. To address this, data augmentation techniques, such as random zooming, brightness adjustments, and horizontal flipping, are employed to artificially expand the dataset and improve model robustness. Additionally, synthetic data generation using the Synthetic Minority Oversampling Technique (SMOTE) ensures a balanced distribution of classes, enhancing the model's ability to generalize across all stages of Alzheimer's.

This research focuses on developing an automated system to classify MRI scans into four distinct categories: Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented. To accomplish this, the study investigates several neural network architectures, including Artificial Neural Networks (ANNs), Convolutional Neural Networks (CNNs), and a transfer learning model based on InceptionV3. Each model is trained using a balanced dataset to assess its performance through metrics such as accuracy, Area Under the Curve (AUC), and loss. The use of data augmentation and the Synthetic Minority Oversampling Technique (SMOTE) is critical for addressing class imbalance and enhancing diagnostic precision.

By evaluating the performance of these models, the study aims to determine the most effective method for Alzheimer's detection. The results contribute to the advancement of computer-aided diagnosis, offering a scalable framework for implementing deep learning models in clinical practice. Ultimately, this research supports early detection initiatives, improving patient outcomes and aiding in the ongoing battle against Alzheimer's Disease.

II. LITERATURE SURVEY

The paper[1] offers an in-depth analysis of machine learning (ML) approaches, with a particular emphasis on Convolutional Neural Networks (CNNs), for the early diagnosis of Alzheimer's Disease (AD). It reviews significant studies conducted between 2016 and 2021 that utilize neuroimaging data, including MRI scans, PET scans, and EEG signals, for classifying AD. The analysis underscores the benefits of deep learning techniques, especially CNNs, in achieving accurate diagnoses, while also addressing their limitations and highlighting potential areas for further refinement.

One of the significant studies mentioned is by Mehmood et al. (2016), who applied a multi-layer CNN to MRI scans for early-stage AD detection. Their approach achieved an impressive accuracy of 98.73%, primarily using tissue segmentation and a deep learning architecture based on the Visual Geometry Group (VGG). The study's strength lies in its high accuracy despite the absence of data augmentation, but it also faces the challenge of being limited by a small sample size of around 300 MRI subjects, which reduces the generalizability of the results.

Another notable study is by Odusami et al. (2021), which used a modified ResNet18 model for feature extraction from structural MRI scans. Their results demonstrated an accuracy of 99.09%, underlining the effectiveness of fine-tuning deep learning models. However, the use of a single dataset, ADNI, raises concerns about the scalability and applicability of the findings across diverse populations. On the other hand, Venugopalan et al. (2020) proposed a 3D CNN model for hippocampus segmentation and analysis, which achieved 78% accuracy in detecting Alzheimer's from MRI scans. While their approach emphasizes reducing noise and enhancing hippocampal features, the model's relatively low accuracy indicates that further refinement of the CNN architecture may be needed for better performance.

Shah et al. (2019) used a hybrid voting classifier method combined with support vector machines (SVM), achieving an accuracy of 84%. This approach focuses on combining multiple classifiers to improve performance, but its lower accuracy compared to CNN-based models points to the need for more advanced techniques in AD detection. Similarly, Islam et al. (2021) applied deep CNNs to analyze brain MRI and PET scans, achieving precision rates of 81% and 82%. The results emphasize the promise of using deep learning techniques for AD diagnosis, though issues such as dataset imbalance and overfitting remain. Thakare et al.

(2018) explored the use of EEG signals combined with SVM for AD detection, yielding a high accuracy of 95%. The advantage of EEG signals lies in their non-invasive nature, making them a potentially cost-effective alternative to other imaging methods, though the lack of spatial resolution in EEG signals compared to MRI scans is a limiting factor.

[1] also compares different types of CNN architectures, such as the 18-layer CNN and the 3D CNN, in the context of AD classification. The 18-layer CNN, with its deeper feature extraction capabilities, outperforms the 3D CNN in terms of accuracy, particularly when working with 2D MRI images. However, while 18-layer CNNs are more effective for 2D imaging, 3D CNNs are better suited for handling volumetric data, such as 3D MRI scans. Despite their advantages, 3D CNNs tend to be computationally more expensive and may not provide as high accuracy as their 2D counterparts when working with the ADNI dataset.

The paper heavily relies on the ADNI dataset, which is widely used in Alzheimer's research due to its comprehensive collection of neuroimaging data and associated clinical information. This dataset provides a solid foundation for model development, but it also has limitations, such as class imbalance, where healthy control cases outnumber AD cases, potentially leading to overfitting. Additionally, the dataset lacks diversity, making it less representative of global populations. Moreover, the ADNI dataset requires extensive pre-processing, which adds complexity to the research and model development process.

A significant challenge highlighted in the paper is class imbalance, which remains a common problem in Alzheimer's research. Many studies encounter the challenge of imbalanced datasets, where the number of healthy control samples significantly exceeds that of Alzheimer's patients. This imbalance can result in biased models that disproportionately favor the majority class. To address this, data preprocessing techniques such as noise reduction, image normalization, and feature extraction are critical for enhancing the quality of medical imaging and improving model performance. Furthermore, the paper touches on the challenge of model interpretability, particularly for deep learning models like CNNs, which are often criticized as "black boxes" in clinical applications. Techniques such as layer-wise relevance propagation (LRP) and guided backpropagation are highlighted as methods to visualize and interpret which brain regions influence the model's predictions.

One of the strengths of the paper is its comprehensive review of various ML techniques applied to early Alzheimer's detection. It provides a clear comparison between CNNs and other models, offering insights into their strengths and weaknesses in the context of neuroimaging data. The paper also emphasizes the practical applications of these ML methods in clinical settings, underlining the potential for early Alzheimer's detection. However, there are some notable limitations. The paper does not delve deeply into the use of data augmentation techniques, which could address the issue of class imbalance and improve model performance, particularly when working with small datasets. Furthermore, while the paper touches on interpretability methods, it does not fully explore the challenges of interpreting deep learning models in the medical field, which is critical for clinical adoption.

Another limitation of the paper is its over-reliance on the ADNI dataset, which, while valuable, may not fully represent the diversity of Alzheimer's patients worldwide. The lack of discussion about the cost of deploying deep learning models in clinical practice is also a gap. The computational resources required for training deep learning models, such as GPUs, can be expensive, and the cost-effectiveness of deploying such models in real-world settings needs further consideration.

In conclusion, the paper offers an insightful and comprehensive review of the current state of machine learning techniques in Alzheimer's disease detection. It highlights the strengths of deep learning methods like CNNs, especially in terms of accuracy, but also identifies key challenges such as class imbalance, data preprocessing, and the need for better interpretability. While the use of the ADNI dataset has provided valuable insights, future research should explore other datasets and address the limitations related to data augmentation, model explainability, and computational cost. By tackling these challenges, machine learning models can become more effective and practical for early Alzheimer's detection in clinical settings.

III. METHODOLOGY

This section details the methodologies used to design and assess deep learning models for classifying Alzheimer's Disease (AD) based on MRI scans. Three distinct approaches were investigated: Artificial Neural Networks (ANNs), Convolutional Neural Networks (CNNs), and a Transfer Learning strategy utilizing the InceptionV3 architecture. Each method focuses on classifying MRI images into four Alzheimer's stages: Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented.

3.1 Dataset used:

The Alzheimer's MRI Dataset, available on Kaggle, is a specialized collection of T1-weighted MRI scans used for the classification of Alzheimer's disease at various stages, from no dementia to moderate dementia. This dataset is organized into four categories corresponding to different stages of the disease:

- Non-Demented (ND) – Individuals showing no signs of dementia.
- Very Mild Demented (VMD) – Early stages of Alzheimer's, characterized by slight cognitive impairment.
- Mild Demented (MD) – Individuals with moderate cognitive decline indicative of Alzheimer's.
- Moderate Demented (MOD) – Advanced stages of Alzheimer's with severe cognitive deterioration.

The dataset is known for its class imbalance, as depicted in Fig3.1, which can present challenges in machine learning tasks. Despite this, it is widely utilized in research to support the development of diagnostic tools aimed at the early detection and classification of Alzheimer's disease.

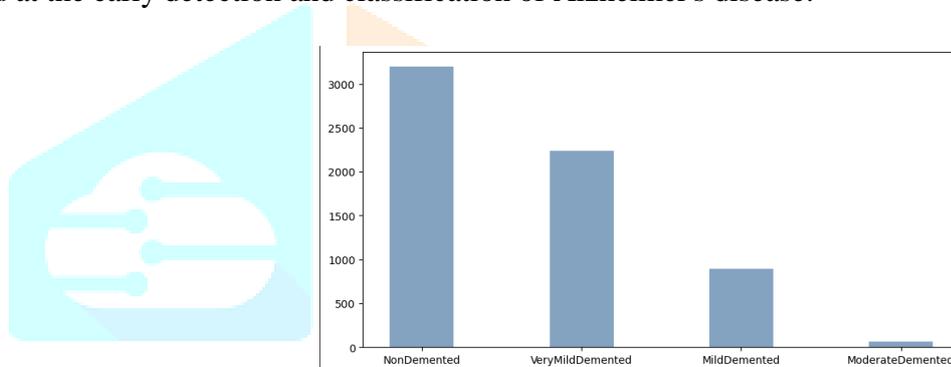


Fig 3.1 Dataset

The MRI images in the dataset have undergone pre-processing steps, including resizing and normalization, to prepare them for model training. These high-resolution scans enable deep learning models, particularly convolutional neural networks (CNNs), to effectively extract features that distinguish between different stages of Alzheimer's disease, enhancing the accuracy of detection and classification models.



Fig 3.2 MildDemented Test Dataset

3.2 Artificial Neural Network (ANN) Methodology

Artificial Neural Networks (ANNs), which are foundational in deep learning, are designed to approximate complex functions by identifying patterns within data. In this study, ANNs were employed as a baseline model to classify the stages of Alzheimer's disease using MRI images.

3.2.1. Architecture Overview

The implemented ANN comprises multiple fully connected (dense) layers. Each layer consists of neurons that take weighted inputs from the preceding layer, apply an activation function to process these inputs, and forward the output to the subsequent layer.

3.2.1.1 Flatten Layer

The ANN begins with a **Flatten layer**. MRI images are 3D arrays of size $176 \times 176 \times 3$. The Flatten layer converts this 3D input into a 1D vector, suitable for input into dense layers.

Input Shape: $(176, 176, 3) \Rightarrow$ Output Shape: $(176 \times 176 \times 3 = 92,928)$

3.2.1.2 Dense Layers

- **First Dense Layer:**

- Contains **100 neurons** with **ReLU** activation.
- The **ReLU** activation function introduces non-linearity into the model, enabling it to learn and capture complex patterns within the data. By replacing negative values with zero while retaining positive values, ReLU helps the network efficiently model non-linear relationships.

- **Subsequent Dense Layers:**

- Four additional dense layers, that contains **200 neurons** and ReLU activation, follow. These layers deepen the network, enhancing its ability to learn intricate data representations.

3.2.1.3 Output Layer

The final layer of the ANN is a dense layer with 4 neurons, corresponding to the four classification categories: Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented. A softmax activation function is applied to this layer, ensuring that the outputs represent probabilities for each class, with their sum equal to 1. This allows the model to predict the most likely class for each input.

3.2.2 Data Preprocessing

3.2.2.1 Image Resizing and Normalization

Images were resized to 176×176 to ensure uniform input size. Each pixel value was normalized to the range $[0, 1]$ by dividing by 255, which helps stabilize training by keeping weights in a similar range.

3.2.2.2 Addressing Class Imbalance with SMOTE

The Alzheimer's dataset showed class imbalance seen in Fig3.1, with certain categories (e.g., Moderate Demented) having fewer samples. **SMOTE (Synthetic Minority Over-sampling Technique)** was applied:

1. **Reshape Input:** Images were reshaped into 1D arrays for SMOTE to operate.
2. **Generate Synthetic Samples:** SMOTE created new samples for minority classes by interpolating between existing ones.
3. **Reshape Back:** The generated data was reshaped back into $176 \times 176 \times 3$.

3.2.3 Model Training

3.1 Loss Function

The model was compiled using Categorical Crossentropy Loss, which is appropriate for multi-class classification tasks. This loss function quantifies the dissimilarity between the predicted probability distribution and the actual class distribution, helping the model adjust its parameters to improve accuracy.

3.2.3.2 Optimizer

The Adam optimizer was employed for its adaptive learning rate capabilities, which effectively combine the advantages of both momentum and RMSProp. This allows for efficient and stable convergence during model training.

3.2.3.3 Early Stopping

To prevent overfitting and promote efficient training, an EarlyStopping callback was implemented. This monitors the model's performance on the validation set and halts training when the specified metric, such as validation loss, stops improving after a set number of epochs.

- Monitors **validation loss**.
- Stops training if validation loss doesn't improve for 3 consecutive epochs.
- Restores best model weights.

3.2.4 Results and Observations

3.2.4.1 Training Process

The model was trained for up to **50 epochs**. During training:

- **Training accuracy** increased steadily, indicating the model's learning capability.
- **Validation accuracy** was monitored to assess generalization.

3.2.4.2 Performance Metrics

The model's performance on the test set was evaluated using:

- **Accuracy:** Proportion of correct predictions.
- **AUC (Area Under the Curve)** measures the model's ability to distinguish between different classes. A higher AUC value indicates better performance in correctly classifying both positive and negative cases, reflecting stronger discrimination between the classes.

3.3 Convolutional Neural Network (CNN) Methodology

Convolutional Neural Networks (CNNs) were utilized to enhance the classification of Alzheimer's disease from MRI images. CNNs are highly effective for image processing tasks because they automatically and efficiently extract spatial and hierarchical features from image data, enabling improved classification accuracy.

3.3.1 CNN Architecture Overview

The CNN implemented in the code is a deep architecture consisting of multiple convolutional, pooling, and fully connected layers. This design is intended to capture both low-level features (such as edges and textures) and high-level features (such as object shapes and patterns) from the MRI images.

3.3.1.1 Input Layer

- **Input Shape:** The input images are resized to $176 \times 176 \times 3$ where:
 - 176×176 represents the image dimensions.
 - 3 corresponds to the RGB color channels.

3.3.1.2 Convolutional Layers

Convolutional layers apply filters to the input data to identify spatial features, such as edges, textures, and shapes.

1. First Convolutional Block:

- **Conv2D:** 16 filters were used along with kernel size 3×3 and ReLU activation.
- **MaxPooling2D:** Reduces spatial dimensions by downsampling, helping to extract dominant features while reducing computational complexity.

2. Second Convolutional Block:

- **Conv2D:** 32 filters were used along with kernel size 2×2 and ReLU activation.
- **MaxPooling2D:** Further downsampling.

3.3.1.3 Separable Convolutional Layers

Separable convolutions are used to reduce computation by splitting the convolution operation into:

1. A depthwise convolution (performs spatial convolution for each input channel separately).
 - A pointwise convolution (using a 1×1 convolution) is applied to combine the outputs from previous layers. This operation helps to mix the features across the channels, allowing the network to capture more complex relationships while maintaining computational efficiency.
- **Third Convolutional Block:**
 - Two **SeparableConv2D** layers with 64 filters each.
 - **Batch Normalization:** Normalizes the activations to stabilize and speed up training.
 - **MaxPooling2D:** Further reduces spatial dimensions.
- **Fourth Convolutional Block:**
 - Two **SeparableConv2D** layers with 128 filters each.
 - **Batch Normalization.**
 - **MaxPooling2D.**

- **Dropout (0.2):** Prevents overfitting by randomly dropping neurons during training.
- **Fifth Convolutional Block:**
- Two **SeparableConv2D** layers with 256 filters each.
- **Batch Normalization.**
- **MaxPooling2D.**
- **Dropout (0.2).**

3.3.1.4 Fully Connected Layers

Once the data passes through the convolutional and pooling layers, it is flattened and then fed into dense layers for classification.

1. **Flatten Layer:** Converts the 3D feature maps into a 1D vector.
2. **Dense Layers:**
 - **512 neurons** used along with ReLU activation, Batch Normalization, Dropout (0.7).
 - **128 neurons** used along with ReLU activation, Batch Normalization, Dropout (0.5).
 - **64 neurons** used along with ReLU activation, Batch Normalization, Dropout (0.3).
3. **Output Layer:**
 - **4 neurons** (one for each class) with **softmax activation** to output probabilities for each class.

3.3.2 Data Preprocessing

3.3.2.1 Image Augmentation

Image augmentation is applied using the **ImageDataGenerator** to enhance the dataset's variability and help the model generalize better:

- **Brightness Adjustment:** Randomly varies brightness.
- **Zoom:** Applies slight zoom-in or zoom-out.
- **Horizontal Flip:** Flips images horizontally.
- **Rescaling:** Normalizes pixel values to the [0,1] range.

3.3.2.2 Class Balancing with SMOTE

To tackle the class imbalance, SMOTE was used to generate synthetic samples for the minority classes. This imbalance is illustrated in Fig3.1.

3.3.3 Model Compilation

3.3.3.1 Loss Function

The model employs Categorical Crossentropy Loss to compute the difference between predicted probabilities and actual labels, making it ideal for multi-class classification tasks.

3.3.3.2 Optimizer

The Adam optimizer is selected for its adaptive learning rate features, which integrate the advantages of both momentum and RMSProp.

3.3.3.3 Metrics

The main evaluation metric is AUC (Area Under the Curve), which measures the model's ability to differentiate between classes..

3.3.4 Training Process

3.3.4.1 Early Stopping

An **Early Stopping** callback is used to track the **validation loss**:

- Stops training when validation loss doesn't improve for 3 consecutive epochs.
- Restores the best weights to prevent overfitting.

3.3.4.2 Training and Validation

The model undergoes training for 10 epochs, with the following results:

- **Training accuracy** increases steadily, indicating that the model is learning patterns.
- **Validation accuracy** helps monitor generalization, ensuring the model performs well on unseen data.

3.3.5 Results and Evaluation

3.3.5.1 Accuracy and AUC

- The model's **accuracy** and **AUC** on the test set provide insight into its classification performance.
- **Loss** values are also tracked to ensure the model is not overfitting.

3.3.5.2 Visualization of Metrics

- **Accuracy Plots:** Show training and validation accuracy over epochs.
- **Loss Plots:** It is used to show the training and validation loss to observe the model's learning dynamics.

3.4 Transfer Learning Using InceptionV3

3.4.1 Model Selection and Preprocessing

The InceptionV3 architecture was chosen for its demonstrated effectiveness in handling complex image classification tasks. The model was pre-trained on the ImageNet dataset, which includes a large variety of natural images. For this study, the pre-trained weights were leveraged to extract high-level features from MRI images.

The dataset consisted of MRI images classified into four categories: Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented. Each image was resized to 176x176 pixels to comply with the input size requirements of the InceptionV3 model. Additionally, the pixel values were normalized to the range [0,1] to maintain numerical stability during training.

3.4.2 Model Architecture Modification

The top layers of the pre-trained InceptionV3 model, originally designed for the ImageNet classification task, were removed. This retained only the convolutional base of the model, which is responsible for extracting features. To adapt InceptionV3 for Alzheimer's disease classification:

- **Global Average Pooling:** A global average pooling layer was added following the convolutional base to reduce the spatial dimensions of the feature maps.
- **Fully Connected Layers:** Several dense layers were added to refine the extracted features. These layers utilized ReLU activations to introduce non-linearity, improving the model's ability to learn complex patterns.
- **Dropout Layers:** Dropout was applied to prevent overfitting, especially given the relatively small size of the training dataset.
- **Output Layer:** The final layer had four neurons, representing the four classes, with a softmax activation function to output class probabilities.

3.4.3 Training Procedure

The training process was divided into two phases: **feature extraction** and **fine-tuning**.

Feature Extraction Phase

In the initial phase, the convolutional base of InceptionV3 was frozen, meaning its weights remained unchanged during training. Only the newly added fully connected layers were trained, enabling the model to learn task-specific patterns while preserving the general image features learned from ImageNet.

Fine-Tuning Phase

After the newly added layers converged, fine-tuning was performed. In this phase, some of the deeper layers of the InceptionV3 base were unfrozen and retrained alongside the fully connected layers. A low learning rate was used during this phase to prevent drastic changes to the pre-trained weights, ensuring that the learned features remained relevant while adapting to the specifics of the Alzheimer's dataset.

3.4.4 Optimization and Regularization

The model was optimized with the Adam optimizer, which adapts learning rates for each parameter. The categorical cross-entropy loss function was employed, as it is ideal for multi-class classification tasks.

To prevent overfitting and improve generalization, the following strategies were implemented:

- **Early Stopping:** Training was stopped if the validation loss did not improve after a specified number of epochs, ensuring that the model kept its best weights.
- **Batch Normalization:** Applied in fully connected layers to stabilize learning and reduce sensitivity to initial weights.

3.4.5 Evaluation Metrics

The model's performance was evaluated on a held-out test set using two key metrics:

- **Accuracy:** Calculate the portion of correctly classified images.
- **Area Under the Curve (AUC):** Reflects the model's capability to differentiate between the various stages of Alzheimer's.

Both metrics provided a detailed idea of the model's classification capability.

3.4.6 Implementation Details

The model was trained for up to 10 epochs in each phase, with early stopping based on validation loss. A batch size of 32 was used to balance memory limitations and training efficiency. The entire training process was carried out on a GPU to speed up computations. Upon completion, the model showed strong generalization, successfully distinguishing between the four stages of Alzheimer's.

IV. RESULT

Table 4.1 displays the results of three different machine learning models—Artificial Neural Network (ANN), Convolutional Neural Network (CNN), and Inception model—evaluating their performance in terms of validation accuracy and validation loss.

1. **Artificial Neural Network (ANN):** The ANN model used achieved a validation accuracy of 0.7466, meaning it correctly predicted approximately 74.66% of the validation data. However, it had a relatively high validation loss of 1.2412, which indicates that the model's predictions had a greater degree of error compared to the other models.
2. **Convolutional Neural Network (CNN):** The CNN model performed better than the ANN, with a validation accuracy of 0.9251, or about 92.51%. It also had a lower validation loss of 0.6826, showing that its predictions were more accurate and had less error than the ANN model.
3. **Inception Model:** The Inception model had a validation accuracy of 0.8947, or roughly 89.47%, which is lower than the CNN but higher than the ANN. Its validation loss was 0.7233, indicating a moderate level of prediction error compared to CNN but still an improvement over the ANN.

Overall, the CNN model achieved the highest validation accuracy and the lowest validation loss, making it the best-performing model among the three for this particular dataset. The Inception model also performed well, with accuracy slightly lower than CNN but better than ANN. The ANN model had the lowest accuracy and the highest loss, indicating it was the least effective model for this task.

Model	Validation Accuracy	Validation Loss
ANN	0.7466	1.2412
CNN	0.9251	0.6826
Inception	0.8947	0.7233

Table 4.1 Accuracy comparison of various models used

Figure 4.1 illustrates the training and validation accuracy (above) and training and validation loss (below) over 50 epochs for the Artificial Neural Network (ANN) model. The above graph shows the model's accuracy, with the green line representing training accuracy and the blue line indicating validation accuracy. Both lines fluctuate considerably, especially in the earlier epochs, indicating some instability in the model's learning process. However, as the epochs progress, the accuracy values generally stabilize, with the training accuracy averaging around 85-90%, while the validation accuracy hovers around 75-85%, suggesting a slight overfitting, where the model performs better on the training data than on the validation data.

The below graph displays the model's loss, where the green line shows training loss and the blue line depicts validation loss. Similar to the accuracy graph, the loss values vary significantly across epochs. While the training loss decreases gradually over time, the validation loss exhibits more volatility, with several peaks indicating moments of increased error. Despite this fluctuation, there is a general trend of decreasing loss, with training loss settling between 0.6 and 1.0, while validation loss occasionally spikes but also tends to decrease, averaging between 0.6 and 1.3. This fluctuation in validation loss further supports the indication of overfitting, as the model struggles to generalize consistently to the validation data. Overall, the ANN model shows potential in learning patterns but may benefit from adjustments to reduce overfitting and improve validation stability.

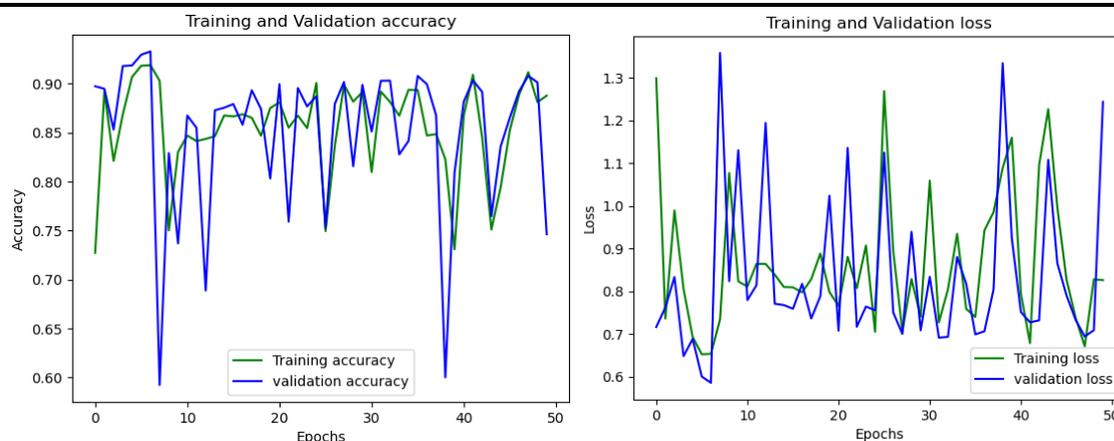


Fig 4.1 Accuracy and Loss for ANN

Figure 4.2 illustrates the training and validation accuracy (top) and the training and validation loss (bottom) for the Convolutional Neural Network (CNN) model over 10 epochs. In the accuracy graph, the green line represents training accuracy, while the blue line indicates validation accuracy. The CNN model demonstrates a rapid improvement in validation accuracy within the initial epochs, rising sharply from around 0.5 to approximately 0.9 by epoch 4. After this initial rise, the validation accuracy stabilizes, showing only minor fluctuations around 90% in the later epochs. The training accuracy, represented by the green line, continues to improve steadily and surpasses 90%, indicating that the CNN model is effectively learning from the training data.

The below graph illustrates the loss values, where the green line represents training loss and the blue line represents validation loss. The training loss shows a consistent downward trend, decreasing smoothly to below 0.5, which suggests that the model is minimizing error on the training data effectively. The validation loss, however, is more volatile. After an initial drop, it spikes significantly around epoch 6 before decreasing again in the later epochs. This volatility in validation loss could indicate some irregularities in the model's generalization to new data, though the overall trend suggests improvement. Despite this fluctuation, the CNN model achieves relatively low validation loss values compared to the training loss, suggesting it has learned meaningful patterns and generalizes reasonably well.

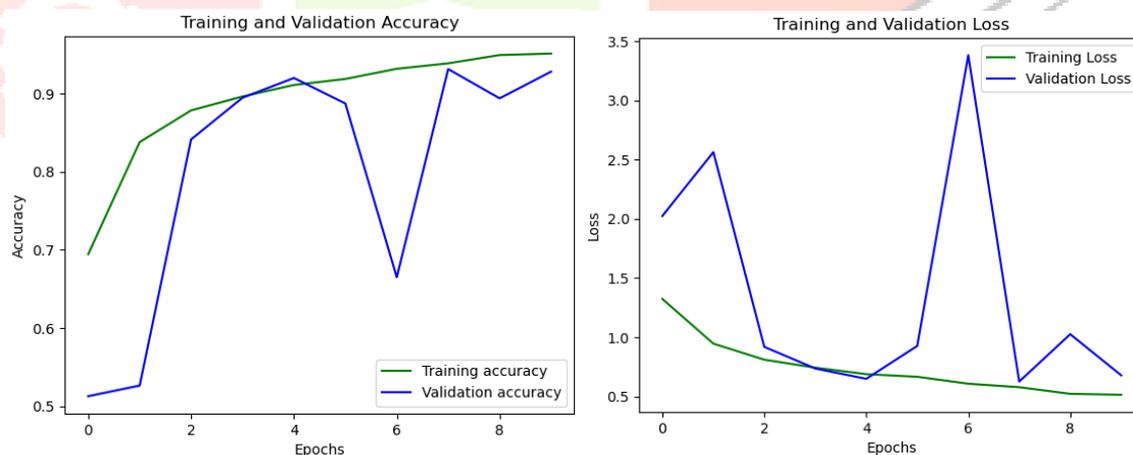


Fig 4.2 Accuracy and Loss for CNN

The Fig4.3 shows two graphs illustrating the training and validation accuracy and loss for an Inception model over multiple epochs. In the accuracy graph (above), the training accuracy (green) steadily improves, reaching around 0.95, while validation accuracy (blue) fluctuates, indicating possible overfitting. In the loss graph (below), training loss (green) consistently decreases, suggesting effective learning. However, validation loss (blue) shows significant spikes, particularly mid-training, reflecting instability and challenges in generalizing to new data. Overall, these patterns suggest that while the model learns well on the training set, further tuning may be needed to enhance validation stability and generalization.

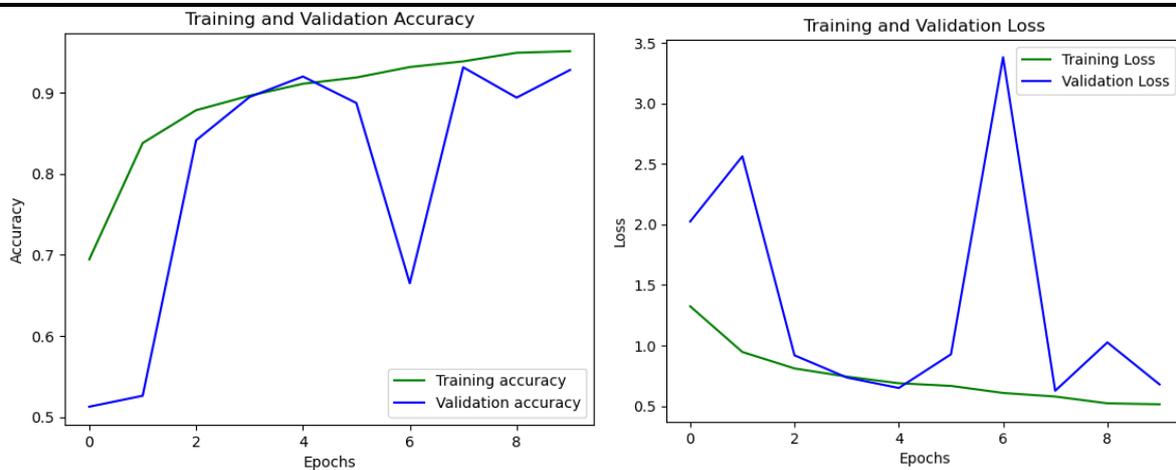


Fig 4.3 Accuracy and Loss for Inception model

V. CONCLUSION

To conclusion, this study compared the performance of three machine learning models—Artificial Neural Network (ANN), Convolutional Neural Network (CNN), and Inception model—for predicting Alzheimer's disease using imaging data. The evaluation metrics, including validation accuracy and validation loss, highlighted significant differences in performance. The CNN outperformed the other models, achieving a validation accuracy of 92.51% and a validation loss of 0.6826, demonstrating its superior predictive capability. The Inception model also performed well, with a validation accuracy of 89.47% and a loss of 0.7233, though it was slightly less effective than the CNN. In contrast, the ANN showed the weakest performance, with a validation accuracy of 74.66% and a higher validation loss of 1.2412, indicating its relative difficulty in making accurate predictions.

These results underscore the significant potential of CNNs in Alzheimer's disease prediction, particularly in tasks that require high accuracy and lower prediction error. While the Inception model showed promising results, the CNN outperformed it, confirming its suitability for this task. The relatively poor performance of the ANN highlights the importance of choosing the right model for complex classification tasks, especially when working with medical imaging data. Future work could focus on further optimizing these models and exploring hybrid approaches to improve accuracy, reduce loss, and enhance the generalizability of the models to real-world Alzheimer's detection scenarios.

REFERENCES

- [1] V. Patil, M. Madgi, and A. Kiran, "Early prediction of Alzheimer's disease using convolutional neural network: A review," *Egyptian Journal of Neurology, Psychiatry and Neurosurgery*, vol. 58, no. 1, p. 130, Nov. 2022, doi: 10.1186/s41983-022-00571-w.
- [2] A. Mehmood, S. Yang, Z. Feng, M. Wang, A. S. Ahmad, R. Khan, et al., "A transfer learning approach for early diagnosis of Alzheimer's disease on MRI images," *Neuroscience*, vol. 460, pp. 43–52, 2021, doi: 10.1016/j.neuroscience.2021.01.002.
- [3] R. Brookmeyer, E. Johnson, and K. Zieglergraham, "Forecasting the global burden of Alzheimer's disease," *Journal of Alzheimer's Disease*, vol. 3, no. 3, pp. 186–191, 2007, doi: 10.1016/j.jalz.2007.04.381.
- [4] World Alzheimer Report 2016. Alzheimer's Disease International. [Online].
- [5] World Health Organization, "Dementia," [Online].
- [6] A. Mehmood, M. Maqsood, M. Bashir, and S. Y. Shuyuan, "A deep siamese convolution neural network for multi-class classification of Alzheimer disease," *Brain Sciences*, vol. 10, no. 2, p. 84, 2020.
- [7] B. X. Bi, S. Li, B. Xiao, Y. Li, G. Wang, and X. Ma, "Computer aided Alzheimer's disease diagnosis by an unsupervised deep learning technology," *Neurocomputing*, vol. 21, pp. 1232–1245, 2019, doi: 10.1016/j.neucom.2018.11.111.
- [8] National Institute on Aging, "What happens to the brain in Alzheimer's disease?" [Online].
- [9] Diogo, V.S., Ferreira, H.A., Prata, D. et al. Early diagnosis of Alzheimer's disease using machine learning: a multi-diagnostic, generalizable approach. *Alz Res Therapy* 14, 107 (2022)