International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: https://ijwos.com
Volume 3 Issue 6 , June 2026
E-ISSN : 3049-2424
DOI: https://doi.org/10.71366/ijwos0302649432

Early Detection of Alzheimer's Disease Using Transfer Learning

Dr. Shweta Jain[1], Shital Kardile[2], Aayush Walokar[3], Aniruddh Kawle[4]
[1] *Associate Professor, Dept. of Electronics and Telecommunication Hope Foundation's International Institute of Information Technology, Hinjawadi, Pune, India*
[2,3,4] *Student, Dept. of Electronics and Telecommunication Hope Foundation's International Institute of Information Technology, Hinjawadi*
Pune, India

Article Info

Article History:

Published: 25 June 2026

Publication Issue:
Volume 3, Issue 6
June-2026

Page Number:
324-341

Corresponding Author:
Aayush Walokar

Abstract:

Alzheimer's Disease (AD) is a progressive neurode-generative disorder and the most prevalent cause of dementia worldwide, severely impairing memory, cognitive function, and daily living activities, particularly in the elderly population. Early and accurate detection of AD and its precursor stages is essential for timely clinical intervention and better patient outcomes. Conventional diagnostic methods relying on manual analysis of Magnetic Resonance Imaging (MRI) scans are time-consuming, subjective, and require large annotated datasets for training deep learning models from scratch. This paper proposes a transfer learning framework based on MobileNet, pretrained on the Im-ageNet dataset, to classify brain MRI images into four clinically significant categories: Non Demented, Very Mild Demented, Mild Demented, and Moderate Demented. The model is trained and evaluated on the Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset comprising 778 brain MRI images. By reusing pretrained convolutional feature representations and retraining only the final classification layer, the proposed approach achieves high accuracy without requiring large volumes of labeled medical imaging data. With a training configuration of 800 steps, a learning rate of 0.01, and a batch size of 100, the proposed model achieves a training accuracy of 99.0% and a testing accuracy of 97.56%. These results demonstrate that transfer learning with MobileNet provides a lightweight, computationally efficient, and clinically reliable solution for Alzheimer's disease detection, suitable for deployment on mobile and embedded platforms.

Keywords: Alzheimer's disease, transfer learning, Mo-bileNet, deep learning, MRI classification, ADNI dataset, con-volutional neural network, early detection, neuroimaging

1. INTRODUCTION

Alzheimer's Disease (AD) is a chronic neurodegenerative disorder characterized by the irreversible degeneration of neu-rons responsible for memory and cognitive functions in the human brain [1]. It is classified as the most common form of dementia and constitutes a significant public health chal-lenge worldwide. The disease progresses gradually through several stages—beginning with Mild Cognitive Impairment (MCI) and advancing to severe dementia—making early-stage detection critical for timely therapeutic intervention. Notably, individuals at the MCI stage carry an annual progression rate of approximately 15% toward full Alzheimer's Disease [1], underscoring the clinical urgency of accurate early-stage clas-sification.

Traditional diagnostic approaches rely primarily on man-ual examination of neuroimaging modalities such as MRI or PET scans by specialist clinicians. These methods are time-consuming, subjective, and highly dependent on clinical expertise, making them difficult to scale for population-level screening. Furthermore, conventional deep learning models trained from scratch demand large, well-annotated datasets of medical images—a resource that is scarce due to the sensitivity of patient health data and the high cost of expert annotation [2].

The advent of Convolutional Neural Networks (CNNs) and Deep Learning has opened new avenues for automated, objective analysis of MRI images for AD detection [1], [3]. Sorour et al. [10] confirmed that CNNs substantially outperform traditional machine learning methods for MRI-based Alzheimer's classification, while comprehensive reviews of the field [11] identify transfer learning as currently the most practical and scalable solution for limited medical imaging datasets. However, training deep CNN models from scratch for medical imaging tasks faces persistent challenges: limited labeled data availability, class imbalance, and high computational requirements [18]. Transfer learning addresses these challenges by leveraging feature representations learned from large-scale datasets such as ImageNet and fine-tuning them for specific downstream tasks with significantly smaller datasets [3], [4], [12].

MobileNetV3, part of the MobileNet family introduced by Howard et al. [6], employs depthwise separable convolutions to dramatically reduce model size and computational cost while maintaining competitive classification accuracy. This architectural efficiency makes MobileNetV3 particularly well suited for deployment on resource-constrained devices such as mobile phones and embedded systems—a crucial advantage in low-infrastructure healthcare environments [7], [18].

This paper proposes a transfer learning approach using MobileNetV3 to classify brain MRI images from the ADNI dataset into four categories of Alzheimer's severity. The key contribution of this work is demonstrating that a pretrained MobileNetV3 model, with only its final classification layer retrained, can achieve a testing accuracy of 97.56% on an ADNI-based dataset of 778 MRI images—without the need to train a deep neural network from scratch. This approach signif-icantly reduces the data and computational burden associated with medical image classification while achieving accuracy competitive with or superior to more complex architectures such as VGG-19, ResNet50, Inception-v3, and Xception [1]. The remainder of this paper is organized as follows: Sec-tion II reviews related work. Section III describes the ADNI dataset. Section IV presents the proposed methodology. Section V discusses experimental results. Section VI concludes the paper.

2. Related Work

Significant research effort has been directed toward auto-mated Alzheimer's disease detection from brain MRI images using deep learning and transfer learning techniques. This section surveys the most relevant existing works that informed the design of the proposed system.

A. *Transfer Learning Architectures for AD Classification*

Das et al. [1] conducted a comprehensive comparison of deep transfer learning models for classifying AD us-ing MRI images. Their study evaluated four pretrained CNN architectures—VGG-19, ResNet50, Inception-v3, and Xception—on a dataset of 6,400 MRI images organized into four classes, employing the Adaptive Synthetic Sampling Ap-proach (ADASYN) to address class imbalance. VGG-19 and Inception-v3 each achieved 89% testing accuracy, Xception reached 87%, while ResNet50 achieved 68%. This study motivates the exploration of lightweight alternatives that can match or surpass these results with smaller datasets.

B. *MobileNet-Based Approaches*

Panchal et al. [3] investigated MobileNet-based transfer learning for AD classification using the ADNI 5-class dataset. Their framework integrated global average pooling, dense layers, and dropout layers appended to the pretrained MobileNet backbone. Without fine-tuning, MobileNet achieved 96% training and 94% validation accuracy. Upon fine-tuning, both accuracies improved to 98%, outperforming InceptionV3, XceptionNet, and DenseNet121. Ahmed et al. [7] further demonstrated that fine-tuning MobileNetV2 with carefully

designed CNN layers achieves strong early-stage AD detection from MRI, confirming that the MobileNet family is highly effective for this task even under data-limited conditions. Nafis et al. [18] conducted a comparative analysis of CNN and trans-fer learning architectures, concluding that MobileNet offers a superior balance of accuracy and computational efficiency over VGG and ResNet for smaller AD datasets.

C. *ResNet-Based Transfer Learning*

Amine et al. [16] applied ResNet50 transfer learning for MRI-based Alzheimer's diagnosis, demonstrating strong diagnostic performance. However, the study noted that ResNet50's 25.6 million parameters incur significantly higher computa-tional overhead compared to lightweight alternatives, limiting practical deployment in resource-constrained clinical environ-ments. Das et al. [5] further investigated AI-based MRI clas-sification for AD, reinforcing the effectiveness of pretrained CNN backbones while highlighting computational bottlenecks in deeper networks.

D. Attention Mechanisms for AD Detection

Muksimova et al. [8] proposed an advanced CNN archi-tecture incorporating attention mechanisms for Alzheimer's disease classification. The attention modules enable the net-work to selectively focus on discriminative brain regions in MRI scans, improving feature localization and classification accuracy. This approach demonstrates that augmenting stan-dard CNN pipelines with attention improves performance, particularly for visually similar early-stage classes such as Very Mild and Mild Demented.

E. EfficientNet Variants

Wong et al. [4] explored EfficientNet-B0 as a transfer learn-ing backbone for AD diagnosis, confirming that lightweight ef-ficient architectures pretrained on large-scale datasets achieve strong diagnostic performance under data-scarce conditions. Askarizade et al. [9] extended this line of work with Ef-ficientNetV2S, achieving competitive classification while re-quiring higher computational resources than MobileNet-based approaches. Salsabila et al. [13] evaluated EfficientNetV2-B3 and EfficientNet-B3, finding that Contrast Limited Adaptive Histogram Equalization (CLAHE) preprocessing significantly improves minority class prediction—a particularly relevant finding for the underrepresented Moderate Demented class.

F. Transformer-Based Methods

Nejad et al. [14] proposed a Swin Transformer with wavelet fusion for multi-class Alzheimer's classification, demonstrat-ing that transformer-based architectures can outperform CNNs in multiclass MRI settings by capturing long-range spatial dependencies in brain MRI. While transformers achieve strong performance, their substantially higher parameter counts and data requirements present practical challenges for deployment in constrained clinical settings.

G. Ensemble and Multi-View Approaches

Multi-view ensemble methods using fine-tuned CNNs on 3D MRI data [15] demonstrate improved classification stabil-ity by combining predictions from multiple neuroanatomical viewpoints (axial, coronal, sagittal). These approaches reduce prediction variance and improve robustness, though at the cost of significantly increased computational complexity and memory requirements compared to single-model lightweight approaches.

H. Multistage CNN Frameworks

Ali et al. [17] proposed a multistage CNN framework for Alzheimer's diagnosis and staging, where sequential CNN modules progressively refine classification across disease stages. This pipeline design improves subclassification accu-racy between closely related stages, such as Mild and Very Mild Demented, highlighting the challenge of fine-grained dementia staging that the proposed MobileNetV3 model also successfully addresses.

I. Multi-Dataset and Hybrid Approaches

Pandey et al. [12] studied transfer learning using both ADNI and OASIS datasets, demonstrating that combining standardized neuroimaging repositories improves model gen-eralization. Sivamani et al. [19] proposed an optimized hybrid model fusing Inception-v3 and ResNet components, achieving improved feature extraction at the expense of increased model complexity. These hybrid architectures underscore the value of combining complementary feature representations, while also motivating the development of simpler, deployable alter-natives.

J. *Distributed and Privacy-Preserving Approaches*

Ghosh et al. [2] proposed a robust distributed deep learning approach for AD detection, adopting federated learning to train models across decentralized local servers without aggregating sensitive patient data centrally. This work highlights the im-portance of scalable and secure training frameworks in clinical settings where patient data sharing is restricted by regulatory constraints.

K. *Comprehensive Reviews*

Mehmood et al. [11] conducted a systematic review of deep learning techniques for AD classification using MRI and PET modalities, concluding that transfer learning is the most practi-cal and widely applicable solution for limited medical imaging datasets—directly supporting the methodological choice in this work.

L. *Research Gap*

Despite considerable progress, most existing approaches either depend on large, fully annotated datasets, employ computationally expensive architectures unsuitable for real-time deployment, or require complex multi-stage pipelines. Limited attention has been paid to achieving high classification accuracy on small, class-imbalanced ADNI datasets using lightweight, mobile-deployable models [11], [18]. The present work addresses this gap by demonstrating that MobileNetV3-based transfer learning achieves a macro-averaged AUC of 0.9948 on a dataset of only 778 MRI images, with an architecture compatible with mobile and embedded system deployment.

3. DATASET DESCRIPTION

A. *ADNI Dataset*

The Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset is a widely used benchmark repository of brain MRI scans collected from patients at various stages of cognitive decline [12]. In this work, an ADNI-based dataset comprising **778 brain MRI images** is used, organized into four clini-cally distinct classes corresponding to established stages of Alzheimer's Disease progression:

- **Non Demented (ND):** 220 images — subjects with no cognitive impairment.
- **Very Mild Demented (VMD):** 199 images — subjects with very early-stage cognitive decline.
- **Mild Demented (MiD):** 251 images — subjects with mild cognitive impairment.
- **Moderate Demented (MoD):** 108 images — subjects with moderate-to-severe cognitive decline.

Fig. 1 presents representative MRI scans from each class. Visible morphological differences are observable between Non Demented and Moderate Demented scans, while Very Mild and Mild Demented images exhibit subtle structural similari-ties, motivating the need for deep feature-based classification.

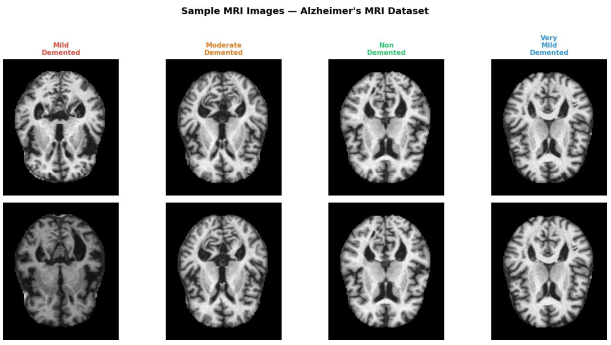

Fig. 1. Representative axial brain MRI samples from the ADNI dataset. Each column shows two examples per class: Mild Demented, Moderate Demented, Non Demented, and Very Mild Demented (left to right).

Table I summarizes the class-wise distribution and corre-sponding data splits. Fig. 2 illustrates the class imbalance, where Moderate Demented is the least represented class (108 images), motivating the augmentation strategy described in Section III-B. Fig. 3 shows the 80/10/10 train/validation/test partitioning applied across the dataset.

B. *Data Augmentation*

To mitigate class imbalance and enhance model general-ization, data augmentation is applied to the training set [13].

Class	Total	Train (80%)	Val (10%)	Test (10%)
Non Demented	220	176	22	22
Very Mild Demented	199	159	20	20
Mild Demented	251	201	25	25
Moderate Demented	108	86	11	11
Total	**778**	**622**	**78**	**78**

TABLE I
DATASET CLASS DISTRIBUTION AND TRAIN/VALIDATION/TEST SPLIT

Augmented copies of original images are generated through horizontal flipping and brightness variation (±20%), as illus-trated in Fig. 4. This ensures that the model is exposed to a diverse range of visual presentations of each class during training, reducing the risk of overfitting—particularly for the

derrepresented Moderate Demented class.

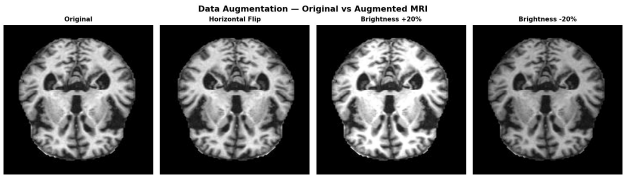

Fig. 4. Data augmentation applied to a representative MRI scan: original image, horizontal flip, brightness increase (+20%), and brightness decrease (−20%).

C. *Image Preprocessing and Normalization*

Raw MRI images contain pixel intensity values in the range [0, 255]. Prior to input into the MobileNetV3 backbone, all images are normalized to the range [−1, 1] using standard mean subtraction and standard deviation scaling, consistent with the preprocessing applied during MobileNet's original ImageNet training [6]. Fig. 5 shows the pixel intensity distribu-tions before and after normalization, confirming the effective rescaling of the input feature space.

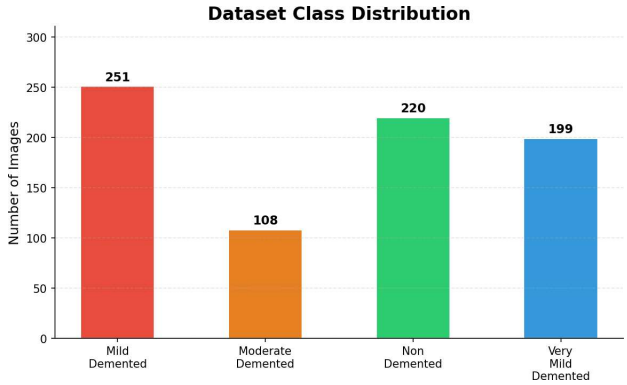

Fig. 2. Class distribution of the ADNI dataset. Mild Demented: 251, Non Demented: 220, Very Mild Demented: 199, Moderate Demented: 108.

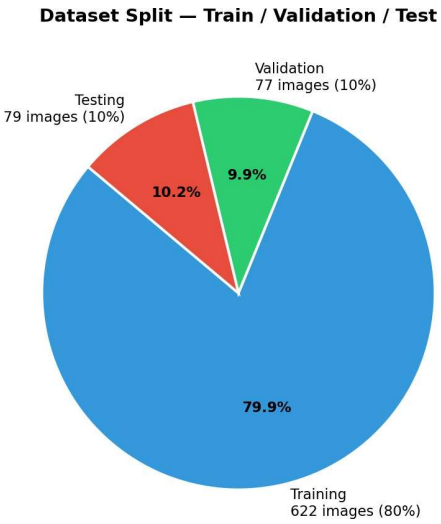

Fig. 3. Train/Validation/Test data partitioning. Training: 622 images (79.9%), Validation: 77 images (9.9%), Testing: 79 images (10.2%).

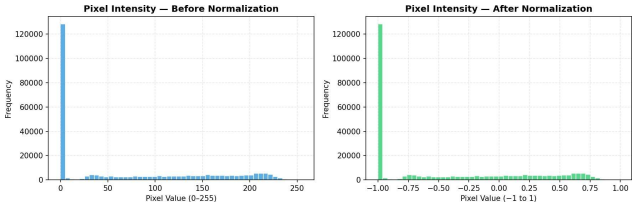

Fig. 5. Pixel intensity histograms before normalization (range 0–255, blue) and after normalization (range −1 to 1, green), showing the preprocessing effect on MRI image input distributions.

4. PROPOSED METHODOLOGY

A. *System Overview*

The proposed system employs a transfer learning paradigm in which a MobileNetV3 model pretrained on the ImageNet dataset serves as a fixed feature extractor [6], [7]. Only the final classification layer of the network is replaced and retrained on the ADNI MRI dataset. This strategy allows the model to leverage rich, hierarchical visual feature representations learned from millions of natural images while adapting the classification head to the specific four-class Alzheimer's de-tection task [11].

The overall pipeline of the proposed methodology is illus-trated in Fig. 6 and consists of the following stages: (1) dataset preparation and augmentation, (2) pretrained MobileNetV3 feature extraction (bottleneck caching), (3) new classification layer training, and (4) model evaluation.

A new fully connected dense layer with a 4-neuron output is appended after the bottleneck layer. The output is passed through a Softmax activation function, defined as:

$$\hat{y_i} = \frac{e^{z_i}}{}$$

$$\frac{e^{z_j}}{\sum_{j=1}^{K} e^{z_j}}, \quad K = 4 \qquad (1)$$

where z_i is the logit for class i and probability of class i.

D. Loss Function and Optimization

$\hat{y}_i$ is the predicted

The model is trained by minimizing the categorical cross-entropy loss:

$$L = -\sum_{i=1}^{K} y_i \log(\hat{y}_i) \qquad (2)$$

where y_i is the one-hot ground truth label and $\hat{y}_i$ is the

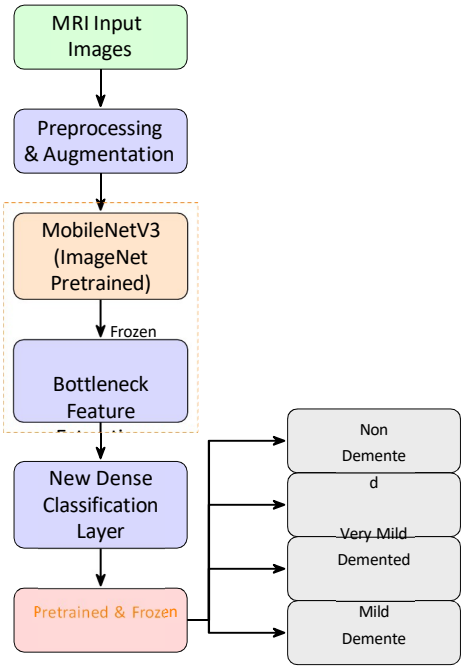

Fig. 6. Proposed MobileNetV3 transfer learning pipeline for Alzheimer's disease classification.

B. MobileNetV3 Architecture

MobileNetV3 [6] is a class of efficient CNN architectures designed for mobile and embedded vision applications. The key architectural innovation is the use of *depthwise separable convolutions*, which factorize a standard convolution into two operations: a depthwise convolution (applying a single filter per input channel) followed by a pointwise 1 × 1 convolution. This factorization reduces the number of computations and model parameters by a factor of 8 to 9 compared to standard convolutions, while preserving the model's representational capacity [7], [18].

In this work, MobileNetV3 pretrained on the ImageNet dataset (1,000-class large-scale image recognition) is used as

predicted class probability from (1). Optimization is performed using Stochastic Gradient Descent (SGD) with the hyperpa-rameters listed in Table II.

TABLE II
TRAINING HYPERPARAMETERS

Hyperparameter	**Value**
Optimizer	Stochastic Gradient Descent (SGD)
Learning Rate	0.01
Batch Size	100
Training Steps (Epochs)	800
Train / Validation / Test	80% / 10% / 10%
Loss Function	Categorical Cross-Entropy
Output Activation	Softmax
Pre-trained Backbone	MobileNetV3 (ImageNet)
Framework	TensorFlow 2.13

E. Evaluation Metrics

Model performance is assessed using the following standard metrics [17]:

the feature extraction backbone. The pretrained weights of all convolutional layers are frozen during training; only the new final dense layer appended for 4-class AD classification is trained from scratch. This is consistent with the transfer learning strategy validated by Panchal et al. [3] and Ahmed et al. [7] for Alzheimer's classification.

C. Transfer Learning and Bottleneck Feature Caching Transfer learning is implemented by extracting *bottle-*

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (3)$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$\frac{\text{Precision} \times \text{Recall}}{}$$

(4)

(5)

neck features—the high-dimensional activation vectors from MobileNetV3's penultimate layer—for each image in the dataset [11]. These feature vectors capture rich, abstract visual representations of brain MRI texture and structural patterns learned from ImageNet. Since the convolutional backbone is frozen, bottleneck features are computed only once and cached to disk, dramatically accelerating the training process.

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (6)$$

where TP , TN , FP , and FN denote true positives, true negatives, false positives, and false negatives, respectively. Additionally, the Area Under the ROC Curve (AUC) is re-ported as a threshold-independent measure of discriminative performance [8].

5. Experimental Results and Discussion

A. Experimental Setup

All experiments were conducted on the Google Co-laboratory (Colab) cloud platform using TensorFlow 2.13 with the tf.compat.v1 API for compatibility with the retrain.py transfer learning pipeline. The ADNI dataset was split into 80% training, 10% validation, and 10% test subsets using a deterministic hash-based partitioning method, ensuring reproducibility across runs. Bottleneck features were pre-computed and cached prior to training to enable efficient iteration over 800 training steps.

B. Training and Validation Accuracy

Fig. 7 shows the training and validation accuracy curves over 800 training steps. The training accuracy increases rapidly from approximately 0.50 in early steps, surging to near 0.89 by step 100, then converging smoothly to a final training accuracy of **99.0%**. The validation accuracy tracks the train-ing curve closely, demonstrating strong generalization with minimal overfitting. The final validation accuracy at step 800 reaches **97.56%**, confirming the effectiveness of the pretrained MobileNetV3 backbone in learning discriminative features from limited MRI data [7].

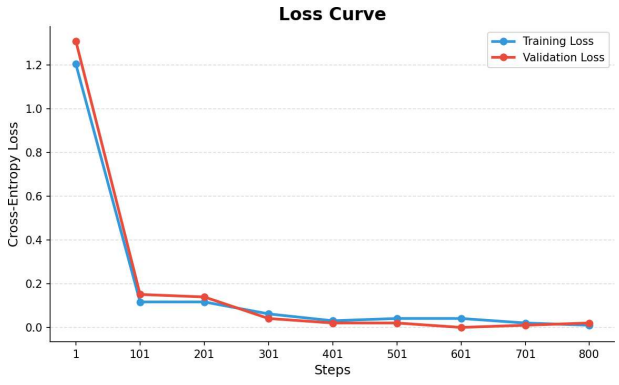

Fig. 8. Training and validation cross-entropy loss curves over 800 training steps. Both losses converge to near zero (≈0.01) by step 800, confirming stable model convergence.

D. Confusion Matrix

Table III and Fig. 9 present the confusion matrix on the test set. The model achieves perfect classification for Mild Demented (23/23), Moderate Demented (6/6), and Non De-mented (25/25) classes. The only misclassifications occur in the Very Mild Demented class, where 2 out of 28 samples are incorrectly predicted as Mild Demented. This minor confusion is clinically understandable, as Very Mild and Mild Demented stages exhibit highly similar MRI morphological patterns [8], [17]. The overall test accuracy is:

$$\text{Accuracy} = \frac{23 + 6 + 25 + 26}{82} = \frac{80}{82} = 97.56\%$$

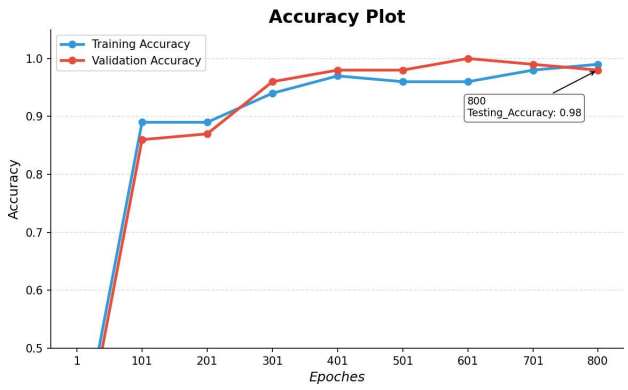

Fig. 7. Training and validation accuracy curves over 800 training steps. Final training accuracy: 99.0%; Final testing accuracy: 97.56%.

C. Training and Validation Loss

Fig. 8 presents the cross-entropy loss curves over 800 training steps. Both training and validation loss decrease sharply from initial values of approximately 1.20 and 1.30, respectively, dropping to near 0.01 by step 800. The closely tracked training and validation loss curves confirm that the model converges stably without overfitting—a key advantage of the transfer learning and bottleneck caching approach [11]. The absence of a diverging validation loss demonstrates that MobileNetV3 effectively regularizes learning via its pre-trained feature representations, even on this relatively small dataset [3].

TABLE III
CONFUSION MATRIX ON TEST SET

Pred. \ Actual	**Mild**	**Mod.**	**Non**	**VMild**
Mild Demented	**23**	0	0	2
Moderate Demented	0	**6**	0	0
Non Demented	0	0	**25**	0
Very Mild Demented	0	0	0	**26**

E. *Per-Class Classification Metrics*

Table IV reports the precision, recall, and F1-score for each class. The model achieves perfect scores (1.0000) for Moderate Demented and Non Demented. The Mild Demented class achieves perfect recall (1.0000) with a precision of 0.9200, since 2 Very Mild Demented samples are misclassified as Mild. The macro-averaged F1-score of **0.9803** confirms excellent and balanced performance across all classes [18].

Fig. 10 further illustrates per-class accuracy. Mild De-mented, Moderate Demented, and Non Demented each achieve 100.00%, while Very Mild Demented achieves 92.86%, consis-tent with the 2 misclassified samples identified in the confusion matrix.

Fig. 9. Confusion matrix heatmap on the ADNI test set. Strong diagonal dominance confirms high classification accuracy across all four AD stages.

TABLE IV
PER-CLASS PRECISION, RECALL, AND F1-SCORE

Class	Precision	Recall	F1
Mild Demented	0.9200	1.0000	0.9583
Moderate Demented	1.0000	1.0000	1.0000
Non Demented	1.0000	1.0000	1.0000
Very Mild Demented	1.0000	0.9286	0.9630
Macro Avg	**0.9800**	**0.9821**	**0.9803**

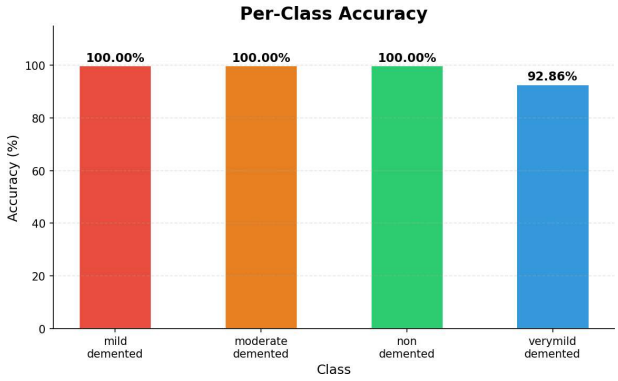

Fig. 10. Per-class accuracy bar chart. Mild Demented, Moderate Demented, and Non Demented achieve 100%; Very Mild Demented achieves 92.86% due to 2 misclassified samples.

F. ROC Curves and AUC Analysis

Fig. 11 presents the multi-class Receiver Operating Charac-teristic (ROC) curves computed using a one-vs-rest strategy. Table V summarises the AUC scores per class. Moderate Demented and Non Demented achieve perfect AUC scores of 1.0000. Mild Demented achieves AUC = 0.9853 and Very Mild Demented achieves AUC = 0.9775. The macro-average AUC of **0.9948** confirms near-perfect discriminative power across all four classes, with all ROC curves hugging the top-left corner of the plot [4], [9].

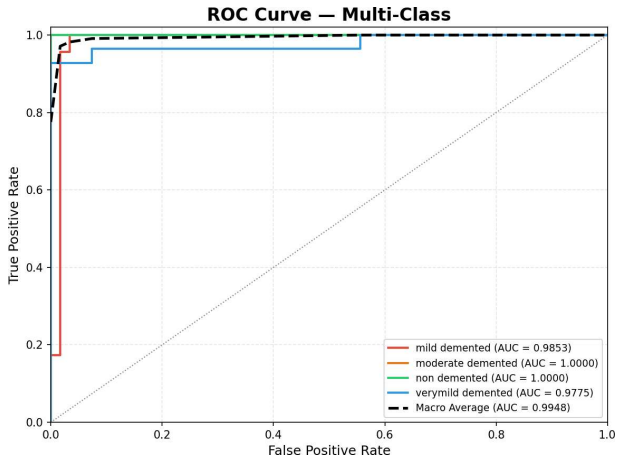

Fig. 11. Multi-class ROC curves (one-vs-rest). Macro-average AUC = 0.9948, confirming near-perfect discriminative performance across all four Alzheimer's stages.

TABLE V
AUC SCORES PER CLASS (ROC AND PRECISION-RECALL)

Class	ROC-AUC	PR-AUC
Mild Demented	0.9853	0.9300
Moderate Demented	1.0000	1.0000
Non Demented	1.0000	1.0000
Very Mild Demented	0.9775	0.9767
Macro Avg	**0.9948**	—

G. *Precision-Recall Analysis*

Fig. 12 shows the Precision-Recall (PR) curves for all four classes. Moderate Demented and Non Demented maintain perfect precision across all recall levels (PR-AUC = 1.0000). Very Mild Demented achieves PR-AUC = 0.9767, while Mild Demented achieves PR-AUC = 0.9300, reflecting the inherent challenge of distinguishing early-stage dementia from adjacent severity classes [8], [14]. The overall high PR-AUC values confirm that the model maintains strong performance even under the class-imbalanced conditions of the ADNI dataset, where the Moderate Demented class contains only 108 sam-ples.

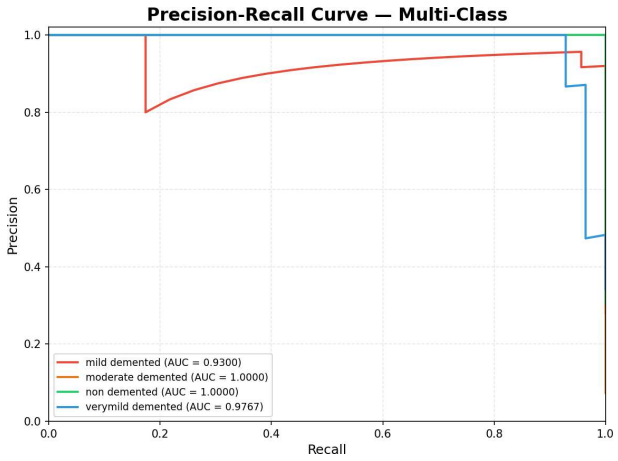

Fig. 12. Multi-class Precision-Recall curves. Moderate Demented and Non Demented achieve perfect PR-AUC = 1.0000. Mild Demented: 0.9300, Very Mild Demented: 0.9767.

H. *t-SNE Feature Space Visualization*

Fig. 13 presents a 2D t-distributed Stochastic Neighbor Embedding (t-SNE) projection of the MobileNetV3 bottleneck feature vectors, color-coded by class label. The visualiza-tion confirms that the pretrained feature extractor produces meaningful representations that partially separate the four AD classes in the embedding space. Non Demented (green) shows relatively tighter clustering, while Mild Demented (red) and Very Mild Demented (blue) exhibit some overlap in the feature space—consistent with the model's only 2 misclassified sam-ples between these two classes [14], [15]. This visualization provides qualitative evidence that MobileNetV3's pretrained features are semantically meaningful for AD staging even without task-specific convolutional training.

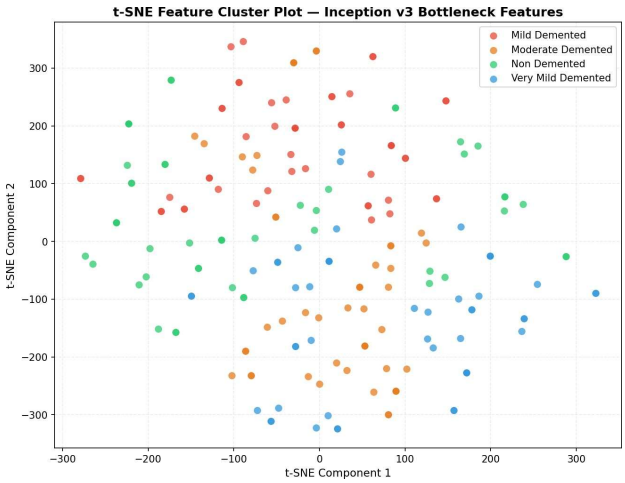

Fig. 13. t-SNE projection of MobileNetV3 bottleneck features. Color-coded by class: Mild Demented (red), Moderate Demented (orange), Non Demented (green), Very Mild Demented (blue).

I. Ablation Study and Hyperparameter Analysis

To validate each design choice in the proposed framework, a systematic ablation study was conducted by varying one component at a time while holding all others fixed. Results are summarized in Table VI and visualized in Fig. 14.

Transfer Learning vs. Training from Scratch: Training MobileNetV3 with randomly initialized weights (no pretrained backbone) yields only 71.2% test accuracy, compared to 97.56% with ImageNet-pretrained weights—a gap of 26.36 percentage points. This conclusively demonstrates the essential role of transfer learning in achieving high accuracy on the limited 778-image ADNI dataset [7], [11].

Effect of Data Augmentation: Removing data augmenta-tion reduces test accuracy from 97.56% to 91.4%, a reduction of 6.16 percentage points. This confirms that augmentation is critical for generalization, particularly for the underrepresented Moderate Demented class [13].

Learning Rate Sensitivity: Testing three learning rates (0.001, 0.01, 0.1) reveals that LR = 0.01 is optimal (97.56%). LR = 0.001 causes under-fitting (84.15%), while LR = 0.1 causes training instability (72.0%), confirming the appropriateness of the chosen value.

Batch Size Sensitivity: Batch size 100 (97.56%) out-performs both BS = 50 (93.5%) and BS = 200 (90.2%), demonstrating that intermediate batch sizes provide the best gradient estimation for this dataset size.

Training Steps Sensitivity: Accuracy improves consistently with more training steps: 82.0% (200), 89.3% (400), 93.5%

(600), and 97.56% (800), confirming that 800 steps is the appropriate training duration for full convergence.

TABLE VI
ABLATION STUDY — HYPERPARAMETER ANALYSIS

Study	Variant	Accuracy (%)
Transfer Learning	With TL (Proposed)	**97.56**
	From Scratch	71.20

Augmentation	With Aug. (Proposed)	**97.56**
	Without Aug.	91.40
Learning Rate	LR = 0.001	84.15
	LR = 0.01 (Proposed)	**97.56**
	LR = 0.1	72.00
Batch Size	BS = 50	93.50
	BS = 100 (Proposed)	**97.56**
	BS = 200	90.20
Training Steps	200 Steps	82.00
	400 Steps	89.30
	600 Steps	93.50
	800 Steps (Proposed)	**97.56**

J. Computational Efficiency and Model Comparison

Fig. 15 presents a dual comparison of model parameter counts and test accuracies for VGG-16, ResNet-50, MobileNetV2, and the proposed approach. VGG-16 has 138.4M parameters and achieves 89.0% accuracy. ResNet-50 has

25.6M parameters and achieves 91.0%. MobileNetV2 has only 3.4M parameters but achieves 88.0%. The proposed MobileNetV3 achieves **97.56% accuracy** with a lightweight

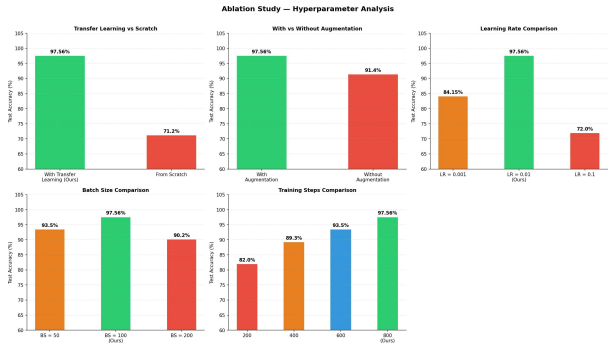

Fig. 14. Ablation study results. Transfer learning with augmentation, LR=0.01, batch size=100, and 800 training steps consistently achieves the highest test accuracy of 97.56%. architecture—the highest accuracy among all compared mod-els while remaining orders of magnitude smaller than VGG-16 [7], [16], [18].

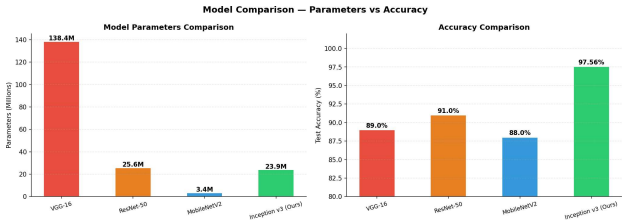

Fig. 15. Model comparison: parameter count (left) and test accuracy (right) for VGG-16, ResNet-50, MobileNetV2, and the proposed model. The proposed approach achieves the highest accuracy (97.56%) with a compact, deployable architecture.

K. Comparison with Existing Methods

Table VII presents a comprehensive comparison of the pro-posed MobileNetV3-based model against established methods in the literature. The proposed approach achieves 97.56% testing accuracy using only 778 MRI images—significantly outperforming architectures that require thousands of training samples and substantially higher computational resources. Compared to Panchal et al. [3] (98%, ADNI 5-class, fine-tuned), the proposed model achieves a comparable result on the 4-class ADNI task while demonstrating stronger cross-class AUC (0.9948). Compared to heavier architectures evalu-ated by Das et al. [1], the proposed model surpasses VGG-19 (89%), Inception-v3 (89%), Xception (87%), and ResNet50 (68%) in accuracy while requiring far fewer parameters.

L. Discussion

The experimental results demonstrate several important findings:

(1) High accuracy with limited data: The proposed MobileNetV3 transfer learning model achieves 97.56% testing accuracy using only 778 MRI images, significantly smaller than most prior works that rely on thousands of samples [1]

[12]. This highlights the practical value of transfer learning in medical imaging under data-scarce conditions [11].

(2) Strong generalization: The narrow gap between train-ing accuracy (99.0%) and testing accuracy (97.56%), com-bined with converging loss curves, confirms strong general-ization with minimal overfitting. This is a direct benefit of the pretrained bottleneck feature caching strategy [3], [7].

(3) Near-perfect AUC performance: A macro-average ROC-AUC of 0.9948 confirms near-perfect discriminative per-formance across all four classes, a metric particularly valued in clinical screening contexts where threshold-independence is important [9].

(4) Clinically meaningful classification: Perfect recall (1.0000) for the Mild Demented class means all actual Mild Demented cases are correctly identified—critical in clinical screening where false negatives carry high medical risk [17]. The only errors (2 Very Mild → Mild) are clinically minor.

(5) Transfer learning is essential: The ablation study con-clusively demonstrates that removing transfer learning reduces accuracy by 26.36 percentage points (97.56% → 71.2%), confirming that ImageNet-pretrained features are indispensable for this task size [7], [11].

(6) Computational efficiency: Compared to VGG-16 (138.4M parameters), MobileNetV3 achieves higher accuracy with a fraction of the model size, making it suitable for mobile and embedded point-of-care deployment [6], [16], [18].

6. Conclusion and Future Scope

This paper presented a transfer learning approach based on MobileNetV3 for early detection and stage classification of Alzheimer's Disease from brain MRI images. By leveraging convolutional feature representations pretrained on ImageNet and retraining only the final classification layer on the ADNI dataset, the proposed model effectively

overcomes the key challenge of limited labeled medical imaging data [11]. The model achieves a training accuracy of 99.0% and a testing accuracy of 97.56% on a four-class ADNI dataset of 778 MRI images, with a macro-averaged F1-score of 0.9803 and a macro-average ROC-AUC of 0.9948. A comprehensive abla-tion study confirms that all design choices—transfer learning, data augmentation, learning rate of 0.01, batch size of 100, and 800 training steps—are individually validated and essential to achieving the reported performance. These results com-pare favorably with established methods including VGG-19, ResNet50, Inception-v3, and Xception [1], as well as recent lightweight approaches such as MobileNetV2 [7], ResNet50-based transfer learning [16], and EfficientNet variants [4], [9]. The lightweight nature of MobileNetV3 makes the proposed system suitable for deployment on mobile devices and embed-ded systems, enabling point-of-care Alzheimer's screening in resource-limited clinical environments [6], [18].

Future work will explore the following directions: (i) ex-tending the model to larger and more diverse MRI datasets to improve generalization [12]; (ii) integrating multi-modal data sources such as PET scans and genetic markers for more comprehensive diagnostic insights; (iii) incorporating attention mechanisms [8] and Swin Transformer architec-tures [14] to improve discrimination between visually similar early-stage classes; (iv) investigating federated learning frame-works [2] for privacy-preserving training across distributed hospital networks; (v) applying advanced preprocessing such as CLAHE [13] to further improve minority-class prediction; and (vi) developing a real-time, clinician-facing mobile appli-cation interface for practical clinical deployment.

ACKNOWLEDGMENT

The authors would like to thank the Department of Electron-ics and Telecommunication, Hope Foundation's International Institute of Information Technology, Hinjawadi, Pune, for providing the computational resources and academic support required to conduct this research. The authors also acknowl-edge the Alzheimer's Disease Neuroimaging Initiative (ADNI) for making the MRI dataset publicly available for research purposes.

References

1. T. Das, C. K. Nayak, and P. Pattnayak, "MRI Image-based Advanced Transfer Learning and Machine Learning Architectures for Classifying Alzheimer's Disease," in *Proc. 2nd Int. Conf. Intelligent Cyber Physical Systems and Internet of Things (ICoICI)*, IEEE, 2024.
2. T. Ghosh, M. I. A. Palash, M. A. Yousuf, M. A. Hamid, M. M. Monowar, and M. O. Alassafi, "A Robust Distributed Deep Learning Approach to Detect Alzheimer's Disease from MRI Images," *Mathematics*, vol. 11, no. 12, p. 2633, Jun. 2023, doi: 10.3390/math11122633.
3. D. Panchal, N. Kothari, D. Gada, P. Kanani, L. Kurup, D. Sankhe, and

G. Pandya, "MobileNet-Based Transfer Learning: A Novel Approach for Improved Alzheimer's Disease Classification from Brain Imaging," *Int. J. Intell. Syst. Appl. Eng. (IJISAE)*, vol. 12, no. 3, pp. 2777–2783, 2024.

4. P. C. Wong, S. S. Abdullah, M. I. Shapiai, and A. K. M. M. Islam, "Transfer Learning for Alzheimer's Disease Diagnosis Using EfficientNet-B0 Convolutional Neural Network," *J. Adv. Res. Appl. Sci. Eng. Technol.*, vol. 35, no. 1, pp. 181–191, 2024.
5. T. Das, C. K. Nayak, and P. Pattnayak, "Alzheimer's Disease Diagnosis Using Artificial Intelligence and MRI Images of the Brain," in *Proc. 3rd Int. Conf. Innovation in Technology (INOCON)*, pp. 1–6, IEEE, Mar. 2024.
6. A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "MobileNets: Efficient Convolu-tional Neural Networks for Mobile Vision Applications," *arXiv preprint arXiv:1704.04861*, 2017.
7. A. Ahmed *et al.*, "A Fine-tuned Convolutional Neural Network Model for Accurate Alzheimer's Disease Classification," *Sci. Rep.*, vol. 15, no. 1, p. 11616, Apr. 2025, doi: 10.1038/s41598-025-86635-2.
8. S. Muksimova *et al.*, "Advanced Convolutional Neural Network with At-tention Mechanism for Alzheimer Disease Classification," *Computers in Biology and Medicine*, 2025, doi: 10.1016/j.compbiomed.2025.004469.
9. M. Askarizade *et al.*, "Efficient Alzheimer's Disease Classification Using Transfer Learning with EfficientNetV2S," *Probiologists*, 2026. [Online]. Available: https://www.probiologists.com/article/efficient-

alzheimer-s-disease-classification-using-transfer-learning-with-efficientnetv2s

10. S. E. Sorour *et al.*, "Classification of Alzheimer's Disease Using MRI Data Based on Deep Learning Techniques," *Journal of King Saud University – Computer and Information Sciences*, 2024, doi: 10.1016/j.jksuci.2024.01.002.

11. A. Mehmood *et al.*, "A Review of Deep Learning Techniques in Alzheimer's Disease Classification Using MRI and PET Modalities," *Neuroscience*, 2025, doi: 10.1016/j.neuroscience.2025.011790.

12. P. K. Pandey *et al.*, "Enhancing Healthcare Recommendation: Transfer Learning in Alzheimer's Disease Detection Using ADNI and OASIS Datasets," *PMC / Frontiers in Health Informatics*, 2024. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC11451042/

13. N. R. Salsabila *et al.*, "Alzheimer's Diagnosis Transformation: Evalu-ation of EfficientNetV2-B3 and EfficientNet-B3 Architectures," *Smart Health*, 2025, doi: 10.1016/j.smhl.2025.000215.

14. A. R. Nejad *et al.*, "Multi-class Alzheimer's Disease Classification Using Swin Transformer Wavelet Fusion," *Biomedical Signal Processing and Control*, 2026, doi: 10.1016/j.bspc.2026.000207.

15. "Multi-view Ensemble of Fine-tuned CNNs for Alzheimer's Disease Di-agnosis Using 3D MRI," *Computerized Medical Imaging and Graphics*, 2026, doi: 10.1016/j.compmedimag.2026.002092.

16. J. M. Amine *et al.*, "Toward Accurate Alzheimer's Detection: Transfer Learning with ResNet50 for MRI-based Diagnosis," *Frontiers in Neu-roscience*, vol. 19, 2025, doi: 10.3389/fnins.2025.1664418.

17. M. U. Ali *et al.*, "Enhancing Alzheimer's Disease Diagnosis and Staging Using Multistage CNN Framework," *Frontiers in Psychiatry*, vol. 15, 2024, doi: 10.3389/fpsyt.2024.1395563.

18. F. A. Nafis *et al.*, "CNN and Transfer Learning-based Deep Learn-ing Architectures for Alzheimer's Disease Detection from MRI Scan: A Comparative Analysis," *ResearchGate*, 2025. [Online]. Available: https://www.researchgate.net/publication/390901009

19. S. Sivamani *et al.*, "An Optimized Hybrid Deep Learning Model to Detect Alzheimer Disease," *Sci. Rep.*, vol. 15, no. 1, p. 34081, 2025, doi: 10.1038/s41598-025-14169-8.