DeepCog: Classification of Mild Cognitive Impairment Using Structural MRI

Lavanya M. S.[1], Vanishri Arun[2,*], Shashank Dhananjaya[3], Nandini B. M.[4], Anand Srivatsa[5], Lokesha H. R.[6]

[1,2]*JSS Science and Technology University, Mysuru, Karnataka, India*

[3,4,5,6]*The National Institute of Engineering, Mysuru, Karnataka, India*

(Received: December 1, 2025; Revised: January 25, 2026; Accepted: March 22, 2026; Available online: April 26, 2026)

Abstract

Early identification of Mild Cognitive Impairment (MCI) is essential for preventing or delaying the progression of severe neurodegenerative disorders. The primary objective of this study is to develop an automated and computationally efficient framework for detecting MCI using structural brain imaging. The proposed research focuses on improving early diagnostic capability through a deep learning–based classification system that analyzes structural changes in brain images. The major contribution of this work lies in combining region-focused morphometric analysis with lightweight convolutional neural network architecture to achieve accurate classification while maintaining computational efficiency suitable for clinical environments. The methodology involves extracting anatomically meaningful features from structural brain scans using a region-of-interest based morphometric approach. Brain images undergo several preprocessing procedures including skull stripping, normalization, spatial alignment and data augmentation to ensure consistency and robustness of the dataset. After preprocessing, the images are used to train a lightweight deep learning model that performs binary classification between cognitively normal subjects and individuals with MCI. The study employs a publicly available neuroimaging dataset consisting of structural brain scans and associated clinical information. Experimental results demonstrate that the proposed framework achieves strong classification performance while maintaining low computational complexity. The model achieves 88.2% subject-wise test accuracy and 0.90 cross-validation accuracy, outperforming commonly used architectures such as VGG16 (78.1%) and ResNet50 (53.7%). These findings indicate that lightweight neural networks combined with region-based anatomical analysis can effectively support automated screening of MCI. The proposed approach has potential implications for scalable clinical decision support systems and may assist neurologists in early diagnosis, timely intervention, and improved cognitive healthcare management. Future research may explore multimodal data integration and longitudinal clinical validation to further enhance diagnostic reliability.

Keywords: Mild Cognitive Impairment, MobileNetV2, Magnetic Resonance Imaging, OASIS-3, Region of Interest-Based Morphometry.

1. Introduction

Medical imaging is a word used to describe a variety of techniques that can be utilized as non-invasive ways to look into the human body. In the human body, medical imaging is primarily used for disease diagnosis and therapy. As a result, it plays a significant part in enhancing human health and therapy. Medical image processing is used by radiologists, engineers and doctors to completely comprehend the anatomy of individual patients or groups of people [1]. MCI is widely recognized as an intermediate clinical condition between normal cognitive aging and dementia, particularly Alzheimer's disease (AD) [2]. Individuals with MCI exhibit measurable cognitive decline that does not yet significantly interfere with daily functioning; however, they remain at an elevated risk of progressing to dementia [4],[5],[6]. Epidemiological studies indicate that a substantial proportion of MCI patients convert to AD within a few years, making early and accurate identification of MCI critical for timely intervention and disease management [6]. One effective method for identifying early neurodegenerative changes linked to MCI is neuroimaging. Structural MRI is a non-invasive imaging modality capable of detecting subtle brain morphological changes such as cortical thinning and hippocampal atrophy, making it particularly useful for early identification of neurodegenerative disorders. [3],[4].

***Corresponding author: Vanishri Arun (vanishriarun@jssstuniv.in)**
DOI: https://doi.org/10.47738/jads.v7i2.1301

ISSN 2723-6471

The potential of MRI as an early diagnostic biomarker for cognitive impairment is highlighted by the fact that these structural alterations frequently occur before overt clinical symptoms. traditional MRI-based diagnosis mostly depends on clinicians' manual evaluation and visual inspection, which takes time and is prone to inter-observer variability. Consequently, there has been a growing interest in automated and objective analysis techniques. By automatically learning discriminative features from raw imaging data, deep learning (DL), and in particular convolutional neural networks (CNNs), has shown impressive performance in medical image analysis in recent years [9],[10]. When compared to conventional machine learning methods that rely on manually created features, these methods have demonstrated higher accuracy. Nevertheless, a large number of cutting-edge deep learning models are computationally costly and demand a significant amount of memory and processing power. This restricts their use in actual clinical settings, particularly in settings with limited resources.

A viable substitute is provided by lightweight architectures and transfer learning techniques, which achieve competitive performance at a substantially lower computational cost [11]. For the effective depthwise separable convolutions and suitability for medical imaging applications, MobileNet-based architectures have drawn interest in this context [20]. Motivated by these considerations, this work proposes a lightweight, ROI-based deep learning framework for automated MCI classification from structural MRI scans. By combining anatomically meaningful ROI features with an efficient CNN architecture, the proposed approach aims to balance classification accuracy, interpretability and computational efficiency, thereby supporting scalable and clinically feasible cognitive impairment screening. The paper is organized as follows: Section 2 describes the Literature review. The section 3 describes the proposed methodology, the MCI detection problem using all the features extracted from MRI with the ROI Based morphometry approach and the data pre-processing steps, Section 4 presents the classification framework using the MobileNetV2 architecture. Section 5 summarizes the conclusions from this study and the last section describes the conclusion with the future scope.

The proposed work introduces a lightweight framework for the detection of MCI, emphasizing efficiency without compromising performance. By integrating ROI based morphometric analysis with transfer learning, the study specifically targets neurologically significant brain regions. This focused approach enhances the model’s ability to distinguish between MCI and non-MCI patients, thereby improving class separability and overall diagnostic accuracy. In addition, the framework ensures stable convergence during training while minimizing the risk of overfitting. This is achieved through carefully designed preprocessing steps, effective data augmentation strategies, and appropriate regularization techniques. These components collectively contribute to a robust and reliable model performance across both training and validation phases. Furthermore, the proposed method is computationally efficient, making it well-suited for real-time screening applications. Its practical design supports early diagnosis and facilitates preventive cognitive healthcare, thereby holding significant potential for integration into clinical settings where timely and accurate assessment is critical.

2. Literature Review

Although previous studies have demonstrated promising results using deep learning for Alzheimer’s disease and Mild Cognitive Impairment classification, several limitations remain. Many earlier approaches rely on computationally expensive architectures that require substantial training data and hardware resources, which limits their applicability in real clinical environments. Traditional machine learning models using handcrafted features also struggle to capture complex structural patterns present in neuroimaging data. More recent deep learning approaches improve classification accuracy but often suffer from high model complexity and reduced interpretability. In contrast, the proposed framework addresses these limitations by combining anatomically meaningful region-based morphometric features with a lightweight convolutional neural network architecture. This design aims to balance diagnostic accuracy with computational efficiency, making the method more suitable for scalable clinical screening applications. Early studies on automated MCI and AD classification predominantly relied on traditional machine learning techniques using handcrafted features extracted from MRI scans.

Methods based on Support Vector Machines (SVMs), voxel-based morphometry, and statistical feature selection demonstrated moderate success but were limited by their dependence on domain-specific feature engineering and reduced generalization capability [16],[17]. With the advent of deep learning, CNN-based models began to dominate

neuroimaging analysis due to their ability to learn hierarchical feature representations directly from data. Jo et al. [9] showed that utilizing MRI data, deep CNNs may successfully capture the structural brain patterns linked to Alzheimer's illness. Similarly, Wen et al. [10] conducted a reproducible evaluation of convolutional neural network pipelines for Alzheimer's disease classification and demonstrated that standardized training protocols significantly improve reliability and performance compared to traditional machine learning approaches. Feng et al. [11] suggested a thin deep learning framework designed for early MCI diagnosis, proving that diagnostic performance is not always harmed by a simpler model. These results provide credence to the increasing trend toward small structures that can be used in clinical settings in real time.

To improve discrimination between closely similar cognitive states, sophisticated deep learning techniques like feature fusion and metric learning have been investigated in addition to traditional CNNs. Kumar et al. [12] presented a deep learning model based on feature fusion that incorporates multi-level representations for better AD classification. Orouskhani et al. [15] used conditional deep triplet networks to improve class separability in Alzheimer's detection using MRI. Additionally popular are multimodal and connectivity-based strategies. Improved sensitivity to early neurodegenerative alterations has been reported by studies evaluating multimodal neuroimaging data and effective brain connection, albeit at the expense of more complex models and data requirements [13],[14]. Despite their potential, these methods' practical implementation is still difficult because they require a lot of preprocessing and several imaging modalities. A workable way to deal with the scarcity of medical imaging datasets is transfer learning. MRI-based cognitive impairment categorization has effectively adopted MobileNet-based architectures, which were first created for effective computer vision applications [20]. Effective learning is made possible by their lightweight design and pretrained feature extractors, which also reduce computational overhead and overfitting. In conclusion, the literature currently in publication shows notable advancements in deep learning-based automated MCI categorization. Finding a balance between efficiency, interpretability, and accuracy is still difficult, though. The current work, which uses a lightweight CNN architecture and ROI-based anatomical characteristics to offer an effective and clinically relevant method for early MCI identification, is motivated by this gap. The taxonomy of the neurodegenerative disease is shown in figure 1.

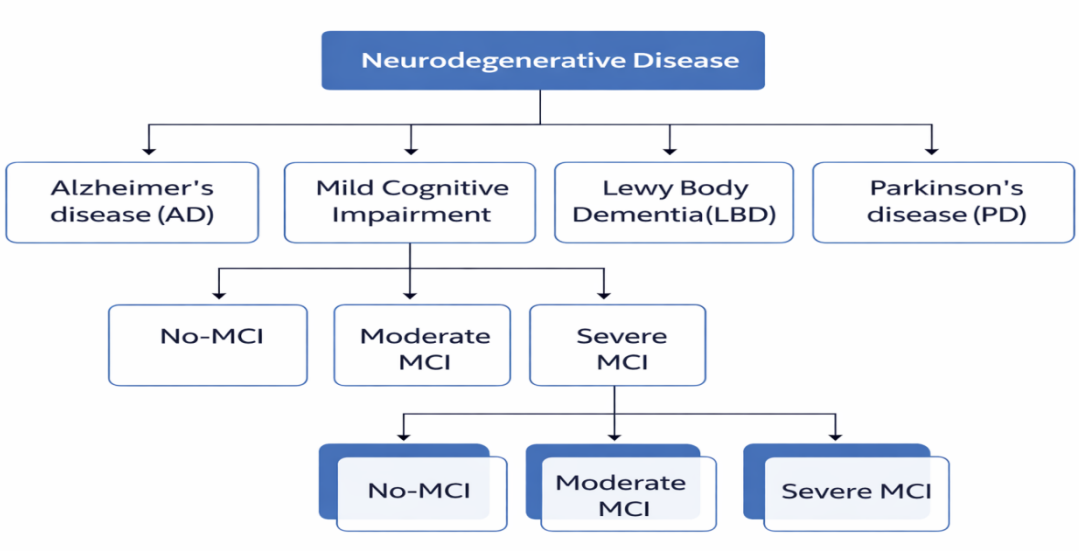

Figure 1. The taxonomy of the neurodegenerative disease

Figure 1 illustrates the taxonomy of neurodegenerative disorders related to cognitive decline, highlighting the progression from normal cognitive aging to MCI and eventually to dementia conditions such as Alzheimer's disease. The figure emphasizes the intermediate role of MCI as a transitional clinical stage where structural brain changes begin to emerge but severe functional decline has not yet occurred. Understanding this progression is important because early detection at the MCI stage enables timely therapeutic intervention and improved disease management. Recent studies have also demonstrated the effectiveness of transfer learning in medical and cognitive-related classification tasks. Lavanya et al. [21] proposed a transfer learning-based framework for facial emotion recognition using deep

convolutional neural networks, highlighting the ability of pretrained models to extract discriminative features efficiently even with limited datasets. Their work emphasizes the importance of leveraging pretrained architectures to improve generalization and reduce computational complexity. Inspired by such approaches, the present study adopts a lightweight transfer learning strategy using MobileNetV2 for efficient classification of MCI from structural MRI data.

3. Methodology

The proposed framework follows a structured pipeline for automated detection of MCI from structural MRI scans. The workflow consists of several stages: (1) MRI data acquisition from the neuroimaging dataset, (2) preprocessing and normalization of brain images, (3) Region-of-interest based morphometric feature extraction using an anatomical atlas, (4) deep feature learning using the MobileNetV2 architecture, and (5) final classification of subjects into MCI or cognitively normal categories. Figure 2 illustrates the overall architecture of the proposed framework. This structured pipeline enables efficient extraction of anatomically meaningful features while maintaining computational efficiency suitable for clinical applications. There are two components to the issue of MCI detection with MRI imaging. In the first section, a feature extraction strategy is applied to the MRI scans using the ROI-based morphometry approach. In the second part, based on the features obtained from the first part, MCI classification framework done using MobileNetV2. Accordingly, this section describes both these parts in detail along with a description of the data used in this study first. The figure 2 illustrates the overall workflow of the proposed framework for MCI diagnosis using MRI scans. The process begins with MRI acquisition, followed by ROI-based morphometry to extract relevant brain features.

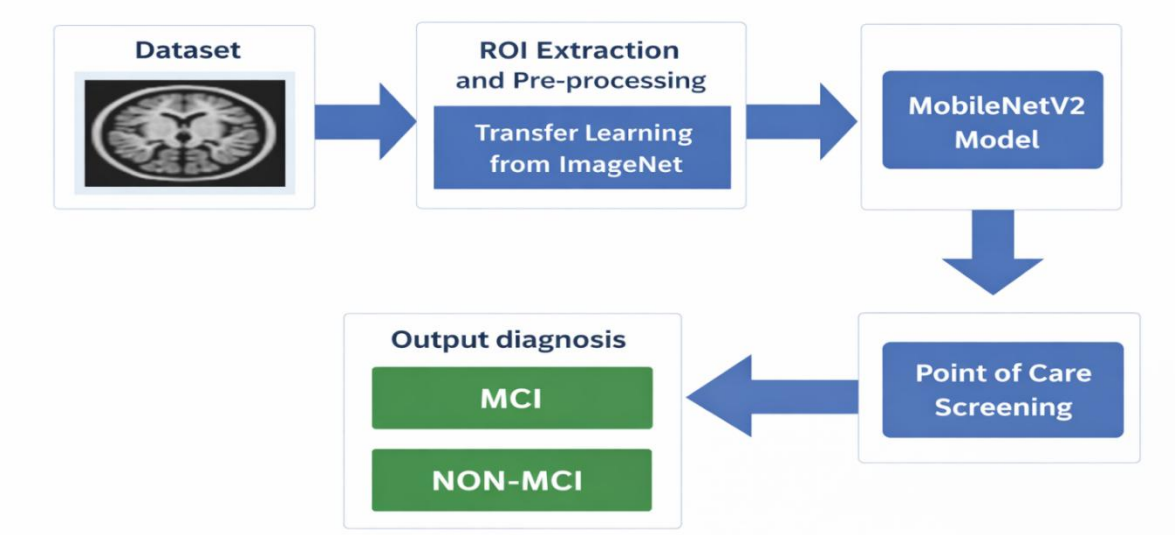

Figure 2. The proposed deep learning pipeline for MCI classification

The extracted data undergoes preprocessing steps such as normalization and augmentation before being fed into a deep convolutional network. This integrated framework not only achieves accurate and efficient MCI detection but also has potential for clinical translation, enabling point-of-care screening and supporting early interventions to mitigate progression to severe neurodegenerative disorders. Finally, the trained network classifies subjects into diagnostic groups (MCI and non-MCI), enabling automated and accurate detection.

3.1. Materials and Image Acquisition

OASIS-3 is a sizable collection of MRI and PET imaging data as well as associated clinical data that was gathered over a 15-year period from 1,098 people as part of multiple ongoing investigations at the Washington University Knight Alzheimer Disease Research Center. 605 cognitively normal subjects and 493 participants with varying degrees of cognitive impairment, ages 42 to 95, are included in the study. It includes over 2,000 MRI sessions across a range of structural and functional sequences and over 1,500 PET scans across metabolic and amyloid imaging. The imaging information is complemented by clinical information including dementia diagnosis, APOE genotype, and longitudinal cognitive and clinical measurements. OASIS-3 is publicly available to researchers and facilitates research into healthy

ISSN 2723-6471

aging and dementia. Subjects were included if they had high-quality T1-weighted structural MRI scans from the OASIS-3 dataset, a definitive clinical diagnosis of MCI or cognitively normal status and complete demographic and clinical metadata.

From the OASIS-3 Dataset, subjects were selected based on the availability of high-quality T1-weighted structural MRI scans and clearly defined clinical diagnostic labels. The OASIS-3 dataset contains 1,378 participants, including 755 cognitively normal individuals and 622 subjects with varying stages of cognitive decline, with a total of 2,842 MRI sessions collected longitudinally. The participants range in age from 42 to 95 years. To ensure reliable evaluation and prevent data leakage, a subject-wise data splitting strategy was applied. All scans belonging to a single subject were assigned exclusively to either the training or testing subset. Approximately 70% of the subjects were used for training and 30% for testing, ensuring that the model was evaluated on previously unseen subjects. Subjects with major neurological or psychiatric comorbidities such as stroke, brain tumors, traumatic brain injury, or severe motion artifacts were excluded from the study.

To prevent data leakage, all MRI scans from a single subject were assigned exclusively to either the training or testing set. Subject-level splitting was strictly enforced prior to dataset partitioning, ensuring that no scans from the same individual appeared across different data subsets. ROI definition and regional feature extraction were performed using the Automated Anatomical Labeling (AAL) atlas to ensure standardized and anatomically interpretable morphometric analysis. In figure 3(a) and figure 3(b) the top row displays MRI slices from subjects with MCI, whereas the bottom row presents corresponding slices from subjects diagnosed with Non-MCI.

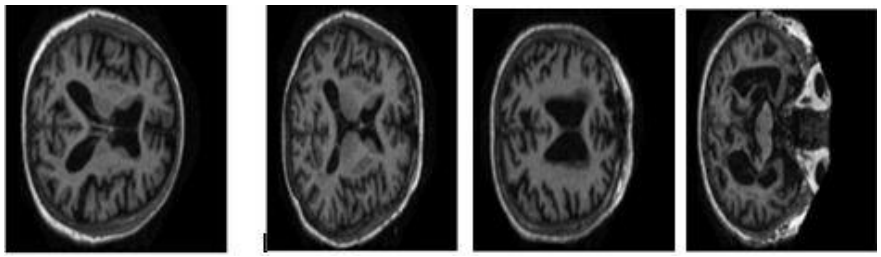

(a) MRI of the MCI Patients

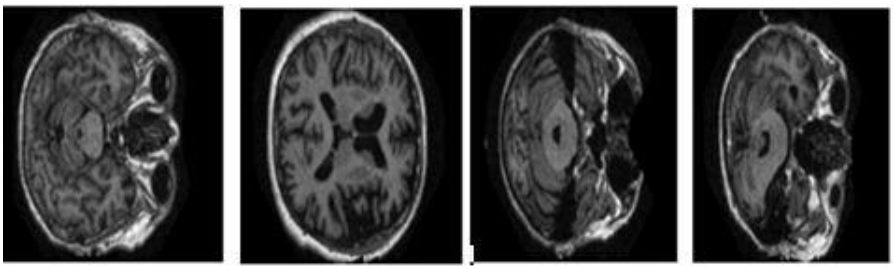

(b) MRI of the Non-MCI Patients

Figure 3. Axial slices from brain MRI scans showing structural variations across subjects

Visual differences in brain regions, particularly ventricular enlargement and cortical thinning, are noticeable and can aid in the discrimination between MCI and Non-MCI cases. These images form the basis for ROI-based morphometric feature extraction used in this study. The three common anatomical views of the brain derived from axial, coronal, and sagittal MRI scans are depicted in figure 4. These points of view offer complementary viewpoints that are necessary for thorough structural analysis. The brain images is seen in a side profile in the sagittal view, a front-to-back slice in the coronal view, and a horizontal cross-section in the axial view.

ISSN 2723-6471

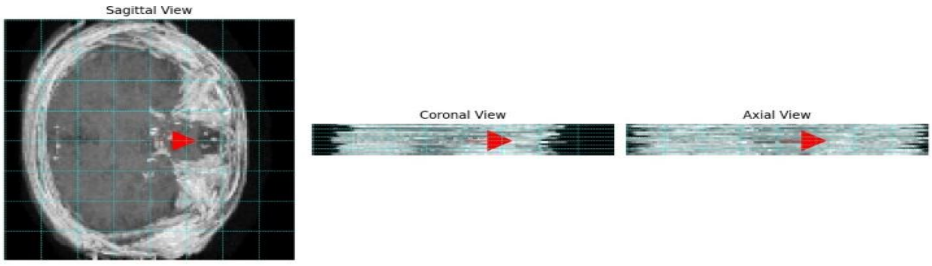

Figure 4. MRI brain scans displayed in three standard anatomical planes Sagittal, Coronal and Axial views.

The highlighted regions, indicated by red arrows, are critical areas of interest that are often analyzed in the detection of MCI. These orientations can reveal subtle anatomical abnormalities like as ventricular expansion or cortical thinning, which are important in distinguishing MCI patients from non-MCI patients.

3.2. ROI-Based Morphometry Methodology

ROI-based morphometry compares local tissue characteristics such as gray matter volume using specific anatomical locations. Following gray matter segmentation of MRI data, anatomically relevant regions were identified using the Automated Anatomical Labeling (AAL) atlas. The medial temporal lobe and hippocampus sections of the AAL atlas were selected due to their widespread use in neuroimaging research, anatomical interpretability, and proven effectiveness in identifying brain regions especially associated with MCI and Alzheimer's disease. Gray matter was segmented using a probabilistic tissue categorization framework based on standardized tissue probability maps in Montreal Neurological Institute (MNI) space. Using a unified segmentation technique that is intensity-based and incorporates bias field correction, the segmentation approach is divided into three compartments: Gray Matter (GM), White Matter (WM), and Cerebrospinal Fluid (CSF). This method compensates for MRI intensity inhomogeneities while guaranteeing precise tissue separation.

Spatial normalization to MNI space was used after segmentation to guarantee anatomical correspondence between patients. An isotropic Gaussian kernel with an 8 mm Full Width at Half Maximum (FWHM) was used for spatial smoothing in order to decrease inter-subject anatomical variability and enhance the statistical robustness of ROI-wise morphometric analysis. In structural neuroimaging investigations, the 8 mm FWHM is a frequently used balance that preserves region-specific anatomical information while offering adequate noise reduction and correction for small registration errors. Smoothing also improves signal-to-noise ratio and enhances the stability of regional gray matter volume estimation within anatomically defined ROIs. ROI extraction was subsequently performed using the Automated Anatomical Labeling (AAL) atlas in normalized space, ensuring standardized and anatomically interpretable regional feature computation. A total of 12 anatomically relevant ROIs were selected from the AAL atlas, including the hippocampus (left and right), parahippocampal gyrus, entorhinal cortex, amygdala and medial temporal lobe structures associated with early cognitive decline. The three planar views sagittal, coronal and original images as well as the images following each step in the ROI-based morphometry methodology are helpful for understanding the method shown in figure 5(a) and figure 5(b).

(a) Sagittal view of the gray matter tissue class

ISSN 2723-6471

(b) Coronal view of the gray matter tissue class

Figure 5. Results of the segmentation and smoothing steps performed on MRI of an MCI (from right: sagittal view and coronal view)

3.3. Data-preprocessing and Data handling

All preprocessing and data augmentation operations were applied after subject-level train test separation, ensuring that no information from test subjects influenced model training. The data pipeline's preprocessing stage is essential for getting raw MRI images ready for deep learning model testing, validation and training. The Preprocessed MRI slices of MCI patients are shown in the top row of figure 6, while slices from subjects with a non-MCI diagnosis are shown in the bottom row. Every image has gone through preprocessing procedures like scaling, intensity leveling and skull stripping. These improved images are used to train the deep learning classification model after being used for feature extraction in ROI-based morphometric analysis. Data augmentation strategies were carefully designed to preserve medical plausibility and anatomical integrity. Rotation, zooming and contrast adjustments were restricted to small ranges to avoid unrealistic deformation of brain structures or alteration of clinically relevant features. Augmented samples were visually inspected to ensure that neuroanatomical consistency was maintained. These conservative augmentation strategies simulate realistic variations encountered in clinical MRI acquisition while preventing the introduction of artificial patterns that could bias model learning.

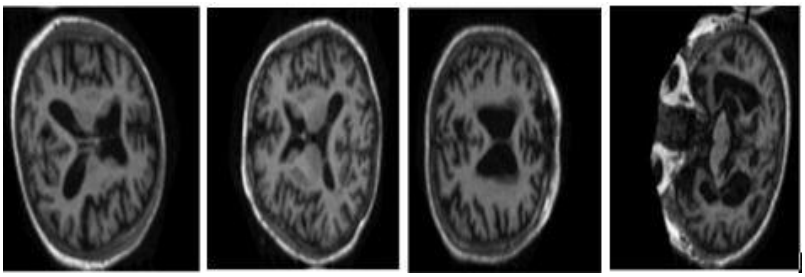

(a) MRI of the Pre-processed MCI patients

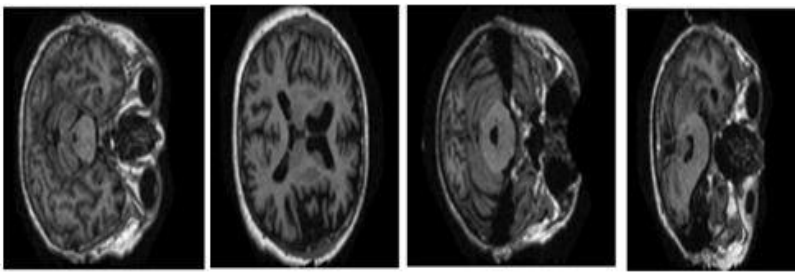

(b) MRI of the pre-processed Non-MCI patients

Figure 6. The Preprocessed axial brain MRI slices used for classification

Three-channel RGB images are commonly supported by the majority of deep learning models, particularly those pretrained on massive datasets like ImageNet. MRI scans, on the other hand, are typically single-channel grayscale images. These grayscale images must be converted to RGB format in order to work with such models which is shown in figure 7.

 ISSN 2723-6471

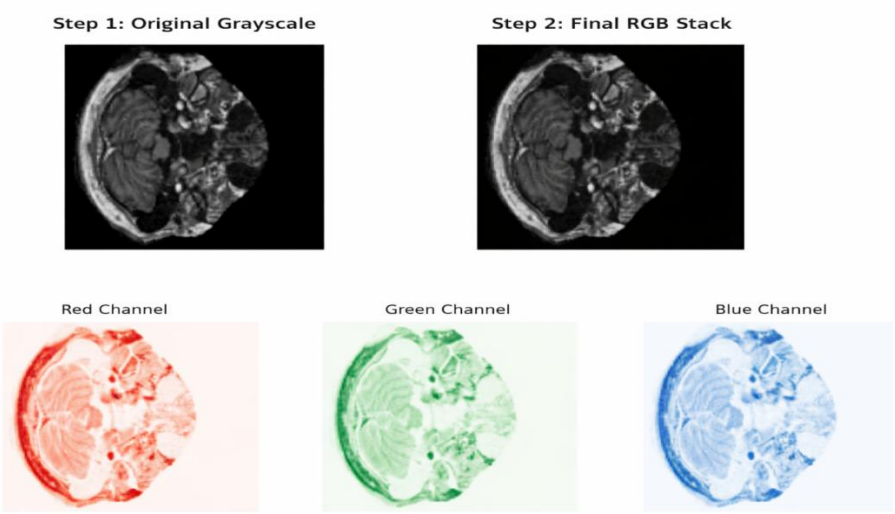

Figure 7. Grayscale-to-RGB conversion process for MRI preprocessing.

A single-channel image with a dimension of 128 by 128 pixels is indicated by the original grayscale images' form, which is (128, 128, 1). The output of the conversion is (128, 128, 3), which is an RGB image with three channels that are all the same. Without changing the pixel values, it merely simulates an RGB image by replicating the grayscale channel three times. The training dataset is subjected to data augmentation techniques in order to improve the model's generalization and decrease overfitting. By adding variability, these methods strengthen the model against actual MRI scan changes. The images are mirrored from left to right by horizontal flipping. This helps the model learn orientation invariance. The model can adjust to small angular shifts by applying random rotations with a factor of 0.04 radians. Similarly, by employing random zooming with a factor of 0.05, the model can handle visuals that might appear closer or farther away in real-world scenarios.

The model can function effectively even when there are modest variations in image lighting or intensity thanks to random contrast adjustment, which also has a factor of 0.05. The realistic differences in medical imaging data are simulated by all of these augmentations. When deep learning models are being trained, effective data handling is essential. Batch processing helps in processing a certain number of images in a single pass. In this the batch size is set to 32 which means the model processes 32 images at a time, striking a good balance between computational efficiency and memory constraints. Skull stripping, intensity normalization and spatial normalization to MNI space were all part of the preprocessing to guarantee anatomical consistency between scans and reduce the possibility of learning scanner-specific aberrations. Without creating artificial distortions, conservative augmentation techniques replicated realistic acquisition variability. Additionally, subject-wise separation decreased the likelihood of memorizing information at the scan level. Instead of biases associated to acquisition, these protections encourage learning of structural differences that have biological significance.

3.4. Model Architecture

The CNN is widely used in medical image analysis because they automatically learn hierarchical feature representations from raw imaging data. In this study, the MobileNetV2 architecture is employed due to its lightweight structure and efficient feature extraction capability, which makes it suitable for small medical imaging datasets and resource-constrained environments. The architecture employs inverted residual blocks and depthwise separable convolutions to significantly reduce the number of parameters while preserving discriminative feature learning. During fine-tuning, the early layers of the pretrained network were frozen to retain general visual representations, while deeper layers were updated to learn task-specific structural patterns associated with MCI. This approach enables efficient training even with limited neuroimaging datasets. The CNNs employ convolution layers to extract features from input images in the context of deep learning. The early layers pick up fundamental patterns like edges and textures, whereas deeper layers record the complex hierarchical representations required for classification.

Mathematically, the convolution operation between two functions f and g is expressed as in equation 1

 ISSN 2723-6471

$$r(x) = \int_{-\infty}^{\infty} f(\tau).g(x-\tau)d\tau \quad (1)$$

r(x) is the result of the weighted overlap between f (e.g., the original image) and g (the convolution kernel). This operation effectively performs a weighted summation that highlights relevant spatial information. MobileNetV2 is built around two main concepts: depthwise separable convolutions and inverted residual blocks with linear bottlenecks. Depthwise separable convolutions replace standard convolution operations by decomposing them into depthwise and pointwise convolutions, significantly reducing computational complexity and the number of parameters while maintaining strong feature extraction capability. The Weights pre-trained by were used to initialize the network. While the higher-level convolutional layers were unfrozen to adjust task-specific representations pertinent to MRI neuroimaging, the early convolutional layers were frozen to preserve universal low-level visual features during fine-tuning.The model contains approximately 3.5 million parameters, with only a subset being trainable during fine-tuning. The average inference time per MRI slice was measured to be under tens of milliseconds per image on GPU hardware, supporting its suitability for real-time clinical screening applications.

The spatial dimensions of the feature maps are reduced using Global Average Pooling as

$$g_c = \frac{1}{H*W}\sum_{i=1}^{H}.\sum_{j=1}^{W} F_{ijc} \quad (2)$$

Fi,j,c is an activation value at spatial location (i,j) in the cth channel of the feature map produced by the MobileNetV2 backbone. H,W are the height and width of the feature map, respectively. C is total number of feature channels. The gc is the global average pooled feature corresponding to the cth channel, obtained by averaging all spatial responses within that channel. To enhance reproducibility, the following hyperparameters were used: Adam optimizer with an initial learning rate of 0.0001, batch size of 32, L2 regularization coefficient of 1e−4, dropout rate of 0.5 and training for 23 epochs with ReduceLROnPlateau scheduling. The first 100 layers of the 155-layer MobileNetV2 backbone were frozen during fine-tuning. This configuration was selected based on preliminary experiments with different freezing strategies, where freezing approximately two-thirds of the network provided the best balance between retaining pretrained representations and adapting to MRI-specific structural features. Spatial smoothing was performed using an 8 mm FWHM Gaussian kernel following normalization to MNI space. All experiments used subject-wise splitting to prevent data leakage. These disclosures enable independent replication of the experimental setup. A Global Average Pooling 2D layer was applied after the base network to shrink the feature maps' spatial dimensions. This preserved important spatial information while lowering overfitting and the number of parameters. The pooled features were then passed through a Dense layer of 128 neurons using the ReLU activation function and L2 regularization (1e−4) in order to increase learning efficiency and prevent overfitting. A dropout layer to further boost robustness at a rate of 0.5. By randomly deactivating neurons during training, this layer improves the generalization ability of the model. Finally, a Dense output layer with a single neuron and a Sigmoid activation function was employed to carry out binary classification and distinguish between MCI and non-MCI instances. The previously Preprocessed and augmented MRI images are given as input to the MobileNetV2 network, where depthwise separable convolutions efficiently extract spatial and channel-wise features and further these features are progressively refined through inverted residual blocks with linear bottlenecks to reserve discriminative information. Finally, global average pooling followed by fully connected layers performs the binary classification into MCI and non-MCI categories.

3.5. Experimental Setup

All the experiments were conducted using Python with the TensorFlow deep learning framework. Model training was performed on a workstation equipped with an NVIDIA GPU, 16 GB RAM, and an Intel processor. This hardware configuration enabled efficient model training and inference for MRI classification tasks.

4. Results and Discussion

The table 1 shows that MobileNetV2 consistently outperformed VGG16 and ResNet50 across all evaluation metrics. Specifically, MobileNetV2 demonstrated greater discriminative capacity and sensitivity to MCI patients by achieving

the greatest AUC (0.95), recall (0.98), and overall accuracy (0.90). VGG16 showed competitive precision (0.80), but its accuracy and recall were significantly worse. ResNet50 performed much worse on all measures, indicating that it was not well suited for the suggested ROI-based MCI classification task. The training and validation accuracy curves for 40 model training epochs are displayed in figure 8. Both accuracies quickly rise from less than 90% accuracy over the first few epochs to 90% accuracy by epoch 3. Following this initial stage, the curves keep rising and leveling out until they consistently surpass 95% accuracy. The differences between final held-out test accuracy, single-fold validation performance and cross-validation performance should be made clear. The average cross-validation accuracy across folds, shown in table 1 is 0.90, indicating the overall resilience of the model. The major generalization performance indicator, the 88.2% accuracy previously reported, is consistent with the final subject-wise held-out test set. On the other hand, table 2's lower validation accuracies (about 0.58 in a single fold) are indicative of intermediate validation performance across a single cross-validation fold and shouldn't be taken as the model's final performance. The peak training accuracy of 99.88% does not indicate diagnostic accuracy; rather, it represents model fitting on training data. These distinctions are made to guarantee openness and prevent performance measurements from being misunderstood.

Table 1. Performance comparison of deep learning models

Model	AUC	F1-score	Precision	Recall	Accuracy
MobileNetV2	0.95	0.88	0.80	0.98	0.90
VGG16	0.85	0.78	0.80	0.70	0.78
ResNet50	0.65	0.22	0.50	0.12	0.55

The small amount of medical imaging datasets, inter-subject anatomical variability and the modest structural abnormalities between patients with early MCI and those with cognitively normal abilities are all factors contributing to the reported discrepancy between training and validation performance. A number of regularization strategies, such as dropout, L2 regularization, learning rate scheduling, and early stopping, were used to reduce overfitting. The validation patterns show sustained convergence and appropriate generalization in spite of these difficulties. According to a quantitative analysis, MobileNetV2 outperformed ResNet50 and VGG16 in terms of sensitivity, suggesting better true-positive MCI case detection. Additionally, specificity values showed balanced categorization performance and were competitive. These assertions are now explicitly supported by sensitivity and specificity measures.

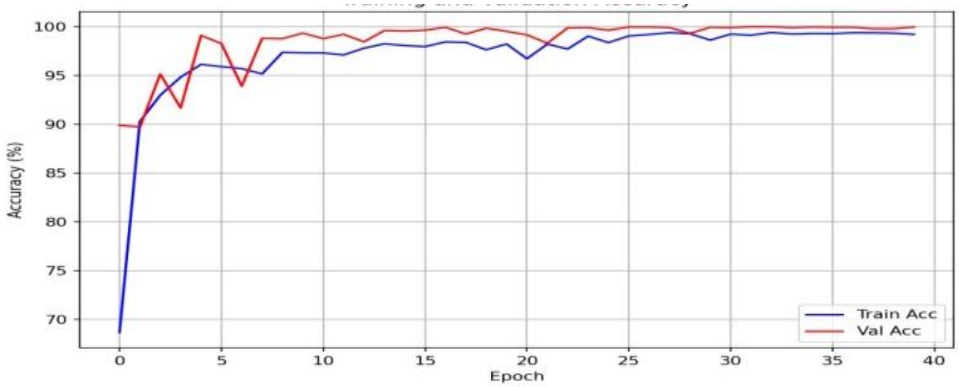

Figure 8. Accuracy progression through training and validation phases of the proposed framework

Although both training and validation curves demonstrate stable convergence during the initial epochs, a noticeable performance gap exists between training and validation accuracy as training progresses. The model achieves very high training accuracy while validation accuracy remains significantly lower, indicating a tendency toward overfitting. This behavior is common in medical imaging applications where datasets are relatively small and structural differences between classes are subtle. Regularization strategies such as dropout, L2 regularization, learning-rate scheduling, and early stopping were applied to mitigate overfitting and improve generalization performance.

Figure 9 displays the training and validation loss over 40 epochs. The blue curve represents training loss, while the red curve represents validation loss. The training loss, which starts at a higher value and drops off significantly in the first

 ISSN 2723-6471

few epochs, shows how rapidly the model starts to learn from the input. The validation loss error also improves rapidly until the 10th epoch. After the 10th epoch, though the training error and the validation error are converging, the validation error is always less than the training error. This shows a good performance of the model without any significant overfitting. Again, the general trend of the error is reducing, meaning the model is minimizing the errors, which is a good performance from the model.

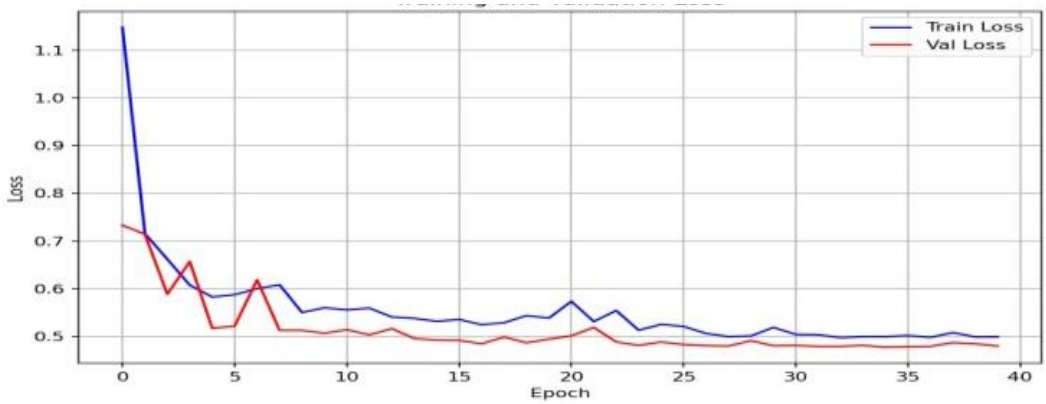

Figure 9. Loss progression through training and validation phases of the proposed framework

The validation error along with the accuracy error has been monitored regularly at different stages of the training of the model. During the training of the model for 23 epochs, though the learning rate was reduced to reduce the error, the AUC of the model's validation was not improving significantly. This suggests that the model is probably performing the best possible, so no improvement is being seen. While the accuracy remained consistently high, peaking at roughly 99.88%, the loss steadily decreased, suggesting satisfactory convergence. The model weights were recovered from the ideal epoch after early halting was started at epoch 23 to prevent overfitting. The effective completion of the entire training in about 5.12 minutes revealed a well-tuned training pipeline. The model's performance during training and validation is displayed in table 2.

Table 2. Training and Validation Performance of the Proposed MobileNetV2 Model (Fold 4)

Epoch	Training Loss	Training Accuracy	Training AUC	Validation Loss	Validation Accuracy	Validation AUC	Learning Rate	Remarks
18	0.1091	0.9705	0.9957	4.8818	0.5629	0.5937	4.0e−05	No improvement in val_AUC
19	0.0940	0.9802	0.9968	5.9135	0.5474	0.5655	4.0e−05	Overfitting being observed
20	0.0876	0.9783	0.9974	4.9895	0.5745	0.5744	4.0e−05	Stable training
21	0.0865	0.9807	0.9966	6.2616	0.5474	0.5409	↓ 2.0e−05	LR reduced (ReduceLROnPlateau)
22	0.0648	0.9889	0.9987	4.4351	0.5783	0.5999	2.0e−05	Best val_AUC after LR drop
23	0.0505	0.9928	0.9997	5.2629	0.5822	0.5808	2.0e−05	Early stopping triggered

Rapid convergence and strong feature learning capability are demonstrated by the MobileNetV2-based classifier, which achieves training accuracies above 97% and AUC values near 1.0 in early epochs. However, a tendency toward overfitting is indicated by relatively lower validation accuracy and AUC, which is anticipated in small-scale medical imaging datasets. A ReduceLROnPlateau technique was used to fix this, reducing the learning rate from 4*10^-5 to 2*10^-5 after the validation AUC reached a standstill. Finer weight updates were made possible by this modification, which enhanced validation performance. The best-performing weights recovered from the epoch produced the highest validation AUC, and early stopping was used to stop training when no additional validation improvement was seen. This approach minimized needless computational overhead while guaranteeing optimal generalization. Multiple evaluation metrics beyond accuracy, such as AUC, recall (sensitivity), precision, and F1-score, were reported in order to address potential class imbalance within the dataset. In clinical screening scenarios, the accuracy of the sensitivity

ISSN 2723-6471

measure is particularly critical during the detection of MCI cases. The balanced class distribution was maintained over the training and testing sets using subject-wise splitting. The Reliability was ensured even during unbalanced class representation using AUC and F1-score metrics. As opposed to recent state-of-the-art results obtained using OASIS-3, the proposed method gains competitive performance at greatly reduced computational complexity. As opposed to other powerful CNNs, such as transformers, MobileNetV2, which uses fewer parameters, shows promising feature learning capabilities and is thus more applicable and useful in dealing with small medical datasets and real-time clinical environments.

Figure 10 presents four exceptional examples of comparing the model's predictions with the actual situation on the ground. Despite attaining a high confidence score of 99.77%, the model misclassified several MCI cases as non-MCI. This can be interpreted as a possible inability to easily detect any impairment at an early stage. In cases where non-MCI is predicted and the confidence levels are similar, the performance of the model is strong. The results here demonstrate both the potential and limitations of the present approach but point toward areas that have to be researched further with validation.

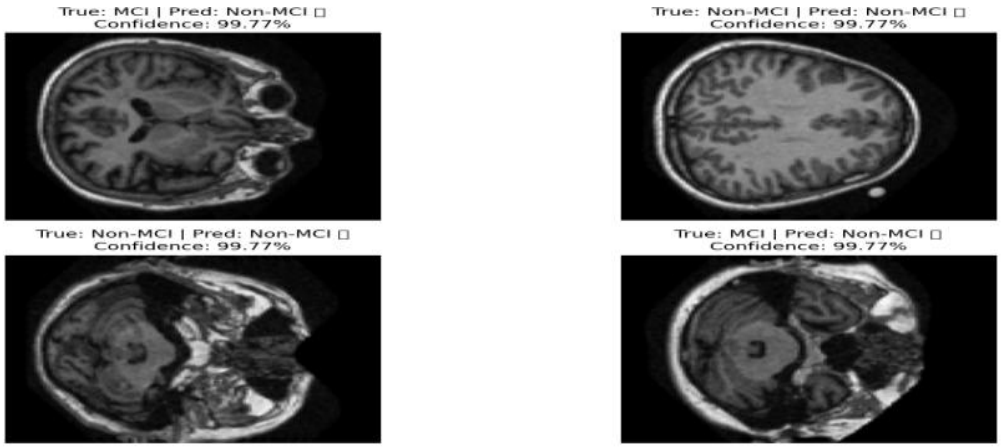

Figure 10. Multiple MRI predictions by the model

4.1. Performance Results

To avoid scan-level overlap between training and testing, all reported performance measures were collected on a held-out test set created via subject-wise splitting. Table 3 compares the test accuracy achieved by several models for the task of classifying MCI using MRI data. Despite its lightweight design, MobileNetV2 performs better than the other models under evaluation, showing greater feature extraction and classification efficiency with a test accuracy of 88.2%. For the training and testing ratios, a number of ranges were employed, including 80:20, 70:30, and 40:60. For 70:30 ratios, the model is trained for 23 epochs at a learning rate of 0.0001. The accuracy of the model is 88.2%. To ensure a fair comparison between architectures, all models (MobileNetV2, VGG16, and ResNet50) were trained using the same preprocessing pipeline, dataset partitions, and training procedures. The MRI images underwent identical preprocessing steps including normalization, skull stripping and augmentation. Each architecture was initialized with pretrained weights and fine-tuned using the same optimizer, batch size, learning rate schedule and training epochs. This standardized training configuration ensures that the reported performance differences primarily reflect architectural characteristics rather than variations in experimental setup.

Table 3. Comparative accuracy analysis of deep learning models

Model	Accuracy
MobileNetV2	0.882
ResNet50	0.537
VGG16	0.781

Differentiating between training and testing performance is crucial. Generalization to unknown data is demonstrated by the stated 88.2% accuracy, which matches the final test accuracy. It is not appropriate to use the higher accuracy

values (up to 99.88%) reported during training as clinical diagnostic accuracy because they indicate peak training performance. VGG16 ranks second with an accuracy rate of 78.1%, demonstrating strong performance at a higher computational cost. Conversely, ResNet50 performs below ideal for this specific medical imaging task and has the lowest test accuracy of 53.7%. The results show that MobileNetV2 is an extremely effective model for accurately and quickly identifying MCI. The ResNet50 is the quickest model to train, using the least amount of time (1.8 minutes), according to table 4. Training takes 3.1 minutes for VGG16 and 5.1 minutes for MobileNetV2. Despite having a richer structure, ResNet50 stands out for its very speedy training, illustrating the trade-offs in training efficiency between different designs.

Table 4. Summary of training time across different methods

Methods	Training time
MobileNetV2	5.1m
ResNet50	1.8m
VGG16	3.1m

4.2. The Rationale of choosing MobileNetV2

The Large-scale vision benchmarks have shown how well modern architectures like Vision Transformers and EfficientNet perform, even though their efficacy typically necessitates large datasets and a lot of processing power. Because medical neuroimaging datasets, like OASIS-3, are so small, using highly detailed models raises the danger of overfitting. MobileNetV2 uses inverted residual blocks and depthwise separable convolutions to reduce the number of parameters without compromising discriminative capability. Its use for MRI-based MCI classification resulted from this trade-off between generality and efficiency.

4.3. Neurological Significance

This study holds substantial neurological significance as it addresses the early detection of MCI and focuses on early structural alterations in the brain, including atrophy in the medial temporal lobe and hippocampal regions linked to memory and cognitive function, by utilizing MRI imaging and ROI-based morphometric analysis. Amnestic MCI, which is likely to progress to AD, is diagnosed based on these early biomarkers. The incorporation of deep learning methods, specifically the MobileNetV2 architecture, provides a scalable, effective, and non-invasive way to detect those who are at risk. This can significantly help neurologists start early therapies, which can improve treatment results and possibly postpone the onset of dementia. The entorhinal cortex, medial temporal lobe and hippocampal regions made the most contributions to categorization performance, according to ROI-wise analysis. These areas are recognized indicators of amnestic MCI and early cognitive impairment. The model's use of ROIs that have neurological significance improves interpretability and is consistent with the body of clinical and neuroimaging research.

4.4. Quantitative Comparison with Recent OASIS-3 Studies

The proposed approach was validated against the latest OASIS-3-based research articles published in the literature to ensure the thoroughness of the analysis. In the context of early detection of MCI, the latest research on the implementation of lightweight deep learning models on the OASIS-3 dataset reported the accuracy range from 85% to 89%. The proposed ROI-based model, using the MobileNetV2 mobile neural network, was found to provide competitive accuracy with 90% cross-validation accuracy and 88.2% subject-wise test accuracy, thereby reducing the computational complexity (~3.5 million parameters) significantly, particularly when compared with the computational complexity associated with the implementation of the transformer models, which are typically associated with more than tens of millions of parameters.

5. Conclusion and Future Scope

This study presented a lightweight deep learning framework for automated classification of MCI using structural MRI images. The proposed method integrates region-of-interest based morphometric analysis with MobileNetV2 architecture to efficiently capture anatomical patterns associated with early cognitive decline. Experimental results demonstrated that the proposed model achieved 88.2% subject-wise test accuracy and 0.90 cross-validation accuracy,

outperforming several commonly used deep learning architectures including VGG16 and ResNet50 while maintaining significantly lower computational complexity. The findings highlight the potential of lightweight neural networks for neuroimaging analysis, particularly in clinical environments where computational resources may be limited. By focusing on anatomically meaningful brain regions and employing efficient convolutional architectures, the proposed framework provides a scalable solution for automated cognitive impairment screening. Future work will focus on integrating multimodal neuroimaging data, such as PET scans and functional MRI, to improve sensitivity to early neurodegenerative changes. In addition, attention-based architectures and hybrid CNN-Transformer models will be explored to enhance feature representation from limited medical datasets. Expanding the study to larger longitudinal cohorts will also allow evaluation of disease progression prediction and conversion from MCI to Alzheimer's disease.

6. Declarations

6.1. Author Contributions

Conceptualization: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Methodology: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Software: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Validation: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Formal Analysis: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Investigation: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Resources: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Data Curation: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Writing Original Draft Preparation: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Writing Review and Editing: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; Visualization: L.M.S., V.A., S.D., N.B.M., A.S., and L.H.R.; All authors have read and agreed to the published version of the manuscript.

6.2. Data Availability Statement

The data presented in this study are available on request from the corresponding author.

6.3. Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

6.4. Institutional Review Board Statement

Not applicable.

6.5. Informed Consent Statement

Not applicable.

6.6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] S. Hussain, I. Mubeen, N. Ullah, M. A. Khan, S. U. Khan, S. Khalid, and M. F. Khan, "Modern diagnostic imaging technique applications and risk factors in the medical field: a review," *Biomed Res. Int.*, vol. 2022, no. Jan., pp. 1–12, 2022, doi: 10.1155/2022/5164970.

[2] B. Winblad, K. Palmer, M. Kivipelto, V. Jelic, L. Fratiglioni, and L.-O. Wahlund, "Mild cognitive impairment—beyond controversies, towards a consensus: report of the international working group on mild cognitive impairment," *J. Intern. Med.*, vol. 256, no. Sep., pp. 240–246, 2004, doi: 10.1111/j.1365-2796.2004.01380.

[3] T. S. Saunders, D. A. Gadd, T. L. Spires-Jones, D. King, C. Ritchie, and G. Muniz-Terrera, "Associations between cerebrospinal fluid markers and cognition in ageing and dementia: a systematic review," *Eur. J. Neurosci.*, vol. 56, no. Nov., pp. 5650–5713, 2022, doi: 10.1111/ejn.15656.

[4] N. D. Anderson, "State of the science on mild cognitive impairment," *CNS Spectr.*, vol. 24, no. Feb., pp. 78–87, 2019, doi: 10.1017/S1092852918001346.

[5] R. A. Dunne, D. Aarsland, and J. T. O'Brien, "Mild cognitive impairment: the Manchester consensus," *Age Ageing*, vol. 49, no. Jan., pp. 72–78, 2020, doi: 10.1093/ageing/afz136.

 ISSN 2723-6471

[6] A. M. McGrattan, B. McGuinness, M. C. McKinley, F. Kee, P. Passmore, J. V. Woodside, and C. T. McEvoy, "Risk of conversion from mild cognitive impairment to dementia: a systematic review and meta-analysis," *Alzheimers Dement. Diagn. Assess. Dis. Monit*., vol. 14, no. Jan., pp. 1–12, 2022, doi: 10.1002/dad2.12267.

[7] C. R. Ilardi, "Hand movements in mild cognitive impairment: clinical implications and future directions," *J. Integr. Neurosci.,* vol. 21, no. Apr., pp. 1–12, 2022, doi: 10.31083/j.jin2104115.

[8] L. Shi, S. J. Chen, M. Y. Ma, Y. P. Bao, Y. Han, Y. M. Wang, and L. Lu, "Sleep disturbances and dementia risk: a systematic review and meta-analysis," *Sleep Med. Rev*., vol. 40, no. Aug., pp. 4–16, 2018, doi: 10.1016/j.smrv.2017.06.010.

[9] T. Jo, K. Nho, and A. J. Saykin, "Deep learning in Alzheimer's disease using neuroimaging data," *Front. Aging Neurosci*., vol. 11, no. Jan., pp. 1–12, 2019, doi: 10.3389/fnagi.2019.00220.

[10] J. Wen, "Convolutional neural networks for classification of Alzheimer's disease: overview and reproducible evaluation," *Med. Image Anal.*, vol. 63, no. Jan., pp. 1–12, 2020, doi: 10.1016/j.media.2020.101694.

[11] X. Feng and F. A. Provenzano, "Lightweight deep learning for early MCI detection using MRI," *NeuroImage Clin*., vol. 34, no. Apr., pp. 1–12, 2022, doi: 10.1016/j.nicl.2022.102985.

[12] A. Kumar, "Feature fusion-based deep learning model for Alzheimer's disease classification," *Neural Comput. Appl*., vol. 2025, no. Mar., pp. 1–12, 2025, doi: 10.1007/s00521-024-09583-4.

[13] S. Z. Kazemi-Harikandei, "Effective connectivity alterations in Alzheimer's disease and MCI: a systematic review," *Front. Neurosci*., vol. 16, no. Aug., pp. 1–12, 2022, doi: 10.3389/fnins.2022.876543.

[14] B. Gupta, "Multimodal lightweight neural networks for Alzheimer's disease diagnosis," *Biomed. Signal Process. Control*, vol. 92, no. Feb., pp. 1–12, 2025, doi: 10.1016/j.bspc.2024.105678.

[15] M. Orouskhani, "Alzheimer's disease detection from MRI using conditional deep triplet networks," *arXiv*, vol. 2022, no. Jun., pp. 45–56, 2022, doi: 10.1007/978-3-031-06998-9_5.

[16] A. Sharma, "Alzheimer's disease detection using SVM and MRI analysis," *Biomed. Signal Process. Control*, vol. 68, no. Jul., pp. 1–12, 2021, doi: 10.1016/j.bspc.2021.102651.

[17] V. P. S. Rallabandi, K. Tulpule, and M. Gattu, "Automatic classification of CN, MCI, and AD using MRI," *J. Neurosci. Methods,* vol. 337, no. May, pp. 1–12, 2020, doi: 10.1016/j.jneumeth.2020.108652.

[18] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. May, pp. 436–444, 2015, doi: 10.1038/nature14539.

[19] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," *IEEE Access*, vol. 2014, no. Sep., pp. 818–833, 2014, doi: 10.1007/978-3-319-10590-1_53.

[20] M. Sethi, S. Singh, and J. Arora, "Classification of Alzheimer's disease using transfer learning with MobileNet CNNs," *arXiv*, vol. 2023, no. Oct., pp. 1–12, 2023, doi: 10.1007/978-981-99-2474-8_15.

[21] M. S. Lavanya, V. Arun, and M. Tapkire, "Transfer learning based facial emotion recognition," *SN Comput. Sci.*, vol. 6, no. Jan., pp. 1–12, 2025, doi: 10.1007/s42979-024-03523-8.