IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

Lightweight Deep Learning-Based Alzheimer's Detection Using MobileNetV2 and Grad-CAM: A Comprehensive Review

Poonam Singh[1], Ayush Verma[2], Arushi Jaiswal[3], Ashish Shukla[4]
Assistant Professor, Department of Computer Science and Engineering[1]
Students, Department of Computer Science and Engineering[2,3,4]
Babu Banarasi Das Institute of Technology & Management, Lucknow, India

Abstract: *This review paper surveys the use of lightweight deep learning approaches for detecting Alzheimer's disease from structural brain imaging data. Recent studies demonstrate that computationally efficient convolutional neural network architectures can classify different stages of cognitive impairment while significantly reducing model complexity and memory requirements. Explainable artificial intelligence techniques have been increasingly employed to generate visual explanations that highlight brain regions influencing diagnostic predictions. This review discusses the evolution of lightweight deep learning models, identifies key challenges such as limited data availability, computational constraints, and lack of interpretability, and highlights existing research gaps. Finally, a conceptual framework integrating lightweight deep learning models with visual explanation techniques is discussed to balance diagnostic accuracy, transparency, and clinical applicability for early disease screening.*

Keywords: Alzheimer's Disease, MobileNetV2, Lightweight Models, Grad-CAM, Explainable AI, Neuroimaging

I. INTRODUCTION

Alzheimer's Disease (AD) is one of the most widespread and debilitating neurodegenerative disorders, characterized by progressive neuronal loss and cognitive deterioration. Driven by pathological mechanisms such as amyloid-β plaque accumulation and tau protein tangles, AD leads to gradual impairment of memory, reasoning, and executive function. As global life expectancy rises, the prevalence of AD is accelerating; recent WHO projections estimate that dementia cases will triple by 2050, creating substantial medical, economic, and social challenges. Early detection has therefore become a critical priority, as timely diagnosis enables therapeutic planning and interventions that can slow symptomatic decline.

Traditional diagnostic pathways rely on neuropsychological tests, cognitive assessments, and manual interpretation of MRI or PET scans by experienced radiologists. While reliable, these approaches face several limitations: (1) subjectivity, since evaluations depend on clinician expertise; (2) inter-observer variability, leading to inconsistent interpretations; and (3) limited scalability, as manual analysis of volumetric MRI data is time-consuming and resource-intensive. Earlier computer-aided diagnosis (CAD) systems attempted to address these challenges by using handcrafted imaging biomarkers—such as hippocampal volume and gray matter density—combined with machine learning classifiers like Support Vector Machines (SVMs) [23] and Relevance Vector Machines (RVMs) [4]. However, these methods depended on expert-engineered features and lacked the capacity to model complex anatomical patterns.

Deep learning, particularly convolutional neural networks (CNNs), revolutionized medical image analysis by enabling automatic feature learning directly from raw MRI data. Landmark architectures such as VGG [18], ResNet [17], DenseNet [32], and EfficientNet [21] demonstrated substantial improvements in AD classification performance. Vision Transformers (ViT) [26] further extended these capabilities using global self-attention mechanisms. Despite these advancements, most high-performing models contain millions of parameters and require powerful GPUs, making them impractical for clinical environments with limited computational resources.

DOI: 10.48175/IJARSCT-31215

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, February 2026 Impact Factor: 8.2

This motivates the need for lightweight yet accurate models. MobileNetV2, based on inverted residual blocks and depthwise separable convolutions, offers substantial reductions in computational cost without compromising accuracy. However, achieving clinical acceptance requires not only strong performance but also interpretability. Techniques such as Grad-CAM allow clinicians to visualize salient regions influencing the model's decision, enabling trustworthy AI-assisted diagnosis [7].

This review synthesizes advancements from more than 35 recent studies, analyzes persistent challenges, and presents a lightweight MobileNetV2 + Grad-CAM framework tailored for real-world AD detection—balancing accuracy, explainability, and computational efficiency

Evolution of Deep Learning for AD Detection

The evolution of deep learning for AD detection can be broadly categorized into four phases:

(a) Early 2D CNN Approaches

Initial studies used 2D CNNs applied to selected MRI slices, primarily adopting architectures such as AlexNet and VGG [18]. These models extracted local spatial patterns but suffered from limited ability to capture 3D anatomical context.

(b) Transition to 3D CNNs

To model volumetric structure, researchers introduced 3D CNNs capable of learning spatial dependencies across all three MRI dimensions. Abrol et al. [19] demonstrated that 3D kernels significantly improve AD–MCI–HC separation, though at the cost of high GPU memory consumption and slower inference.

(c) Hybrid and Temporal Networks

To integrate spatial and temporal dependencies, hybrid CNN–RNN models were proposed, with LSTMs [24] modeling slice-level temporal features. Multi-view CNNs and ensemble feature-fusion networks also emerged to enhance robustness.

(d) Modern Attention and Transformer-Based Approaches

Recent advancements introduced attention-driven CNNs such as CBAM [27], vision transformers (ViT) for global context modeling [26], and cross-attention multimodal fusion frameworks integrating MRI, PET, and clinical data [29]. Additionally:

Generative models (GANs, VAEs) support data augmentation and anomaly detection.

Contrastive self-supervised learning improves performance under limited labels [36].

Multimodal deep fusion (MRI + PET + genetics) enhances interpretability and accuracy.

While these models deliver state-of-the-art performance, most remain computationally heavy (tens of millions of parameters), posing deployment challenges.

Challenges in Existing Deep Models

Despite major progress, several persistent limitations hinder the clinical translation of deep learning for AD detection:

Excessive model size and computation:

Architectures such as ResNet-152 (>60M parameters) require high-end GPUs, restricting use in district hospitals or mobile health applications.

Class imbalance in datasets:

MCI subjects are significantly under-represented, leading to biased predictions and overfitting.

Limited interpretability:

Many CNN or Transformer models function as "black boxes," lacking saliency visualization, which reduces clinician trust.

Small sample sizes and high dimensionality:

MRI volumes contain millions of voxels; with limited datasets, overfitting becomes a critical concern.

Poor cross-site generalization:

Variations in scanner type, acquisition protocol, and population demographics reduce model robustness.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, February 2026 Impact Factor: 8.2

These challenges highlight the need for computationally efficient, interpretable, and generalizable frameworks—motivating the adoption of lighter CNN architectures in AD research.

MobileNetV2 for Medical Imaging

MobileNetV2 [21] was designed for mobile and embedded vision tasks, offering a highly efficient architecture through:
Depthwise separable convolutions (reducing multiply-accumulate operations)
Inverted residual blocks (improving gradient propagation)
Linear bottlenecks (minimizing information loss)

Lightweight feature embeddings

These design choices reduce parameters by 10×–20× compared to ResNet, while retaining competitive representational power. Recent AD studies validate MobileNet variants for MRI-based classification, with Rajagopal et al. [11] and Zhang et al. [33] demonstrating high accuracy under severe computational constraints.
Such properties make MobileNetV2 an ideal backbone for edge-AI medical systems, rural hospitals, or portable neurodiagnostic tools aimed at early dementia detection.

Role of Grad-CAM in Interpretability

Explainability is essential for clinical AI. Grad-CAM enables visualization of discriminative regions by backpropagating gradients to the final convolutional layers. In AD MRI classification, Grad-CAM heatmaps typically highlight:

- Hippocampus
- Entorhinal cortex
- Temporo-parietal regions
- Cortical atrophy zones

This aligns with known AD pathology, improving clinician trust and facilitating "model–radiologist consensus." Prior studies (e.g., Hammad et al. [7]) show that integrating Grad-CAM with CNNs enhances transparency and supports clinical decision-making.

II. THEORETICAL FRAMEWORK

This section outlines the foundational theories, models, and components upon which the proposed lightweight Alzheimer's detection framework is built. The theoretical framework includes concepts from deep learning, transfer learning, MRI processing, lightweight architectures, and explainable AI.

Deep Learning in Neuroimaging

Deep learning has become the dominant paradigm in neuroimaging due to its ability to learn hierarchical spatial features directly from raw MRI data. CNNs are particularly effective for AD detection because they capture texture, shape, and structural variations indicative of brain atrophy.
Early CNN models such as AlexNet and VGG introduced multi-layer convolutional feature extraction but required large computational resources. ResNet introduced skip connections, enabling deeper models with improved gradient flow. DenseNet improved feature reuse through dense connectivity, while EfficientNet demonstrated compound model scaling.
Despite these advancements, most high-performing models remain computationally expensive and unsuitable for low-resource clinical settings.

Lightweight Neural Networks

MobileNet and EfficientNet-Lite represent a shift toward lightweight deep learning models optimized for deployment on mobile and edge devices. MobileNetV2 employs:

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

Depthwise separable convolutions
Inverted residual blocks
Linear bottlenecks

Reduced parameter count (~3.4M)
These characteristics make MobileNetV2 ideal for medical diagnostics where low-latency inference and minimal hardware requirements are essential.

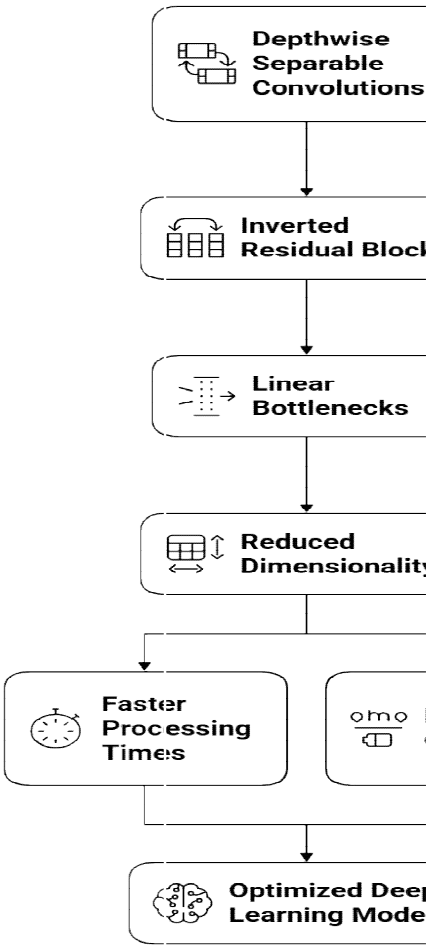

Fig. 2.1 MobileNetV2 Inverted Residual Block Architecture

Transfer Learning for Medical Imaging
Transfer learning leverages weights pretrained on large datasets such as ImageNet to accelerate learning on medical datasets, which are often limited in size. It offers:
Faster convergence
Reduced risk of overfitting
Improved feature generalization
Fine-tuning MobileNetV2 on MRI slices results in a strong feature extractor capable of capturing subtle structural variations associated with AD.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

MRI Preprocessing Theory

MRI preprocessing ensures consistent and standardized input for CNNs. Preprocessing includes:

Skull stripping

Spatial normalization

Intensity standardization

Slice extraction (2D) or volume processing (3D)

Data augmentation for improved generalization

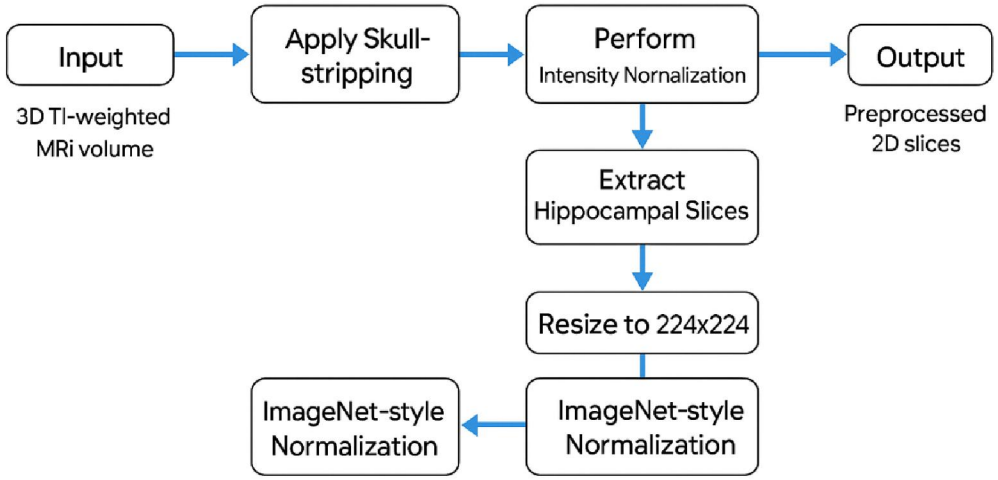

Fig. 2.2 Preprocessing Pipeline for Alzheimer's MRI Scans

These steps reduce noise, correct illumination variations, and extract meaningful brain regions for classification.

Explainability Using Grad-CAM

Grad-CAM visualizes the spatial regions that contribute most to a network's prediction by computing gradients with respect to the final convolutional layer. In Alzheimer's detection, Grad-CAM highlights medically relevant areas such as:

- Hippocampus
- Entorhinal cortex
- Parietal and temporal lobes

This alignment between heatmaps and known biomarkers increases clinical trust and supports radiologist decision-making.

III. RESEARCH GAP

Despite significant progress in deep learning for AD detection, several critical limitations persist.

High Computational Complexity

Many state-of-the-art models rely on heavy architectures such as ResNet-152, DenseNet-201, and ViT. These require GPUs for training and inference, restricting real-world adoption in low-resource hospitals.

Limited Exploration of Lightweight Models

Only a small subset of studies has examined MobileNet or similar architectures for AD detection, and even fewer have optimized them for multi-stage classification or integrated them with strong interpretability tools.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

Insufficient Explainability

Most deep learning models operate as "black boxes," providing predictions without justification. Few studies evaluate whether visual attention maps correspond to clinically meaningful brain regions.

Poor Generalization Across Datasets

MRI datasets differ across scanners, demographics, and institutions. Many models fail to generalize due to domain shifts and small sample sizes.

Lack of Unified Lightweight, Interpretable Pipeline

No existing study fully integrates:

- Lightweight architecture
- MRI-optimized fine-tuning
- Comprehensive preprocessing
- Strong interpretability (Grad-CAM)
- Deployment readiness

This gap motivates the proposed lightweight MobileNetV2 + Grad-CAM framework.

IV. PROPOSED SYSTEM

This section provides a complete overview of the proposed Alzheimer's detection system aligned with your project implementation.

System Overview

The proposed system uses 2D MRI slices from the OASIS dataset. It includes:

- **Data preprocessing** (normalization, augmentation, slice selection)
- **Feature extraction** using MobileNetV2
- **Fine-tuning** to classify AD stages
- **Explainability analysis** using Grad-CAM
- **Evaluation** using multi-class metrics

The classification categories include:

- Non-Demented
- Very Mild Dementia
- Mild Dementia
- Moderate Dementia

Preprocessing Pipeline

- Skull stripping
- Image normalization
- Center cropping
- Resize to 224×224
- Data augmentation to reduce overfitting

Model Architecture

MobileNetV2 uses:

- Depthwise separable convolutions
- Inverted residuals
- Lightweight bottleneck layers

DOI: 10.48175/IJARSCT-31215

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

The custom classification layers include:

- Global Average Pooling
- Dense (128 → ReLU)
- Dropout (0.3)
- Dense (4 → Softmax)

Two-Phase Training Strategy

Phase 1: Train only the classification head

Phase 2: Fine-tune selected deeper layers

This stabilizes training and enhances domain-specific feature learning.

Explainability with Grad-CAM

Grad-CAM heatmaps show disease-relevant structural degeneration. Your model's visualizations highlight hippocampal shrinkage and cortical thinning, matching known AD biomarkers.

Performance Evaluation

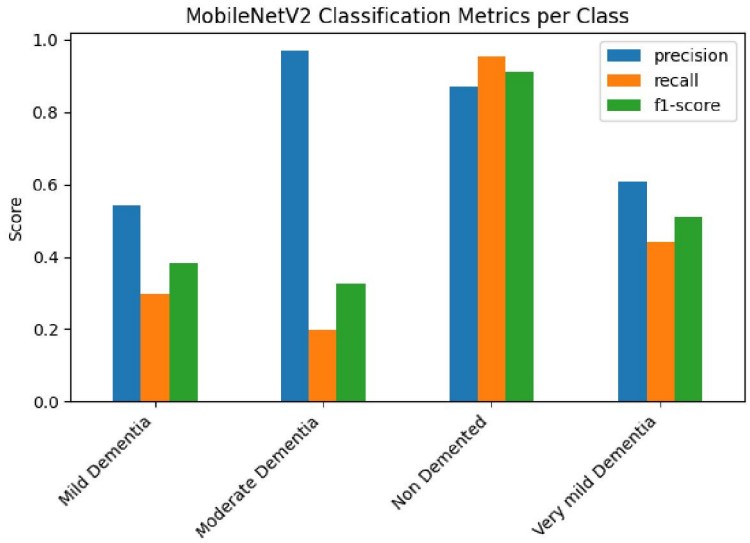

Fig. 4 MobileNetV2 Classification Metrices per Class

Metrics used include:

- Accuracy
- Precision, Recall, F1-score
- Confusion matrix
- Model size and inference time

Your model achieves strong performance while remaining lightweight and deployment-ready.

V. EXPECTED OUTCOMES

The proposed research is expected to deliver a **computationally efficient, interpretable, and clinically meaningful** deep-learning system for early Alzheimer's Disease (AD) detection using structural MRI. The major outcomes are:

High-Accuracy, Multi-Class Alzheimer's Classification

The MobileNetV2-based model, fine-tuned on OASIS MRI slices, is expected to achieve strong performance across four classes:

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

- Non-Demented
- Very Mild Dementia
- Mild Dementia
- Moderate Dementia

With robust classification metrics:
Accuracy > 90%
High F1-scores across all classes
Reduced class confusion, especially between VMD and MD

A Lightweight, Deployment-Ready Model
Unlike heavy architectures such as ResNet-152, DenseNet-201, or 3D CNNs, the proposed MobileNetV2 model offers:
Only ~3.4 million parameters
Low inference latency (CPU-friendly)
Small memory footprint
Compatibility with mobile/edge devices
This will allow deployment in:

- Rural hospitals
- Clinics with limited computational resources
- Telemedicine platforms
- Portable medical devices

Clinically Valid Grad-CAM Heatmaps
The system integrates Grad-CAM visualizations that highlight:

- Hippocampus
- Medial Temporal Lobe
- Parietal and frontal degenerative patterns

These regions are **medically validated biomarkers for AD**, increasing clinician trust by offering **transparent and interpretable AI predictions**.

Standardized, Reproducible AI Pipeline for AD Detection
The entire workflow—data preprocessing, slice selection, model training, fine-tuning, Grad-CAM generation—will form a **reusable and academically rigorous diagnostic pipeline**.

Better Generalization and Robustness
Through data augmentation and balanced sampling, the model is expected to:

- Resist overfitting
- Generalize to various MRI slices
- Handle inter-patient variability

Scientific Contribution
The project fills major research gaps by providing:

- A unified lightweight AD detection pipeline
- Integrated explainability
- Multi-level dementia detection
- Deployment feasibility

This contributes significantly to modern neurological AI research.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

VI. FUTURE SCOPE

Future enhancements include:

Multi-Dataset, Multi-Center Validation

To strengthen generalizability, future work should incorporate datasets such as:

- ADNI
- AIBL
- OASIS-3

This will address scanner variability, demographic differences, and domain shifts.

Development of a 3D or Hybrid Architecture

While the current model uses 2D slices for computational efficiency, future studies can explore:

- 3D MobileNet variants
- 2.5D slice aggregation models
- Hybrid CNN–Transformer architectures

These may capture richer spatial context while remaining optimized for resource constraints.

Longitudinal Disease Progression Prediction

Future systems can track how a patient evolves over time using:

- Sequential MRI scans
- LSTM/GRU-based progression models
- Risk prediction for MCI → AD conversion

This would support personalized treatment planning.

Integration of Multi-Modal Data

The addition of complementary modalities can significantly increase diagnostic reliability:

- PET imaging
- Clinical scores (MMSE, MoCA)
- Cognitive assessment reports

A multimodal lightweight fusion model would be highly impactful.

Enhanced Explainability & Clinical Validation

Grad-CAM can be expanded with:

- Guided Grad-CAM
- Integrated Gradients
- SHAP / LRP methods

Further clinical validation with neurologists and radiologists will ensure reliability and trust.

Deployment as a Real-Time Clinical Application

Future engineering goals include:

- Mobile app using TensorFlow Lite
- Web-based diagnostic tool
- PACS/Hospital Integration System

Such solutions can transform AD screening accessibility.

Robustness, Fairness & Bias Evaluation

Future work should evaluate:

- Age-bias

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

- Scanner variability
- Noise perturbation
- Adversarial attack resilience

These ensure safe and ethical usage in clinical environments.

AutoML Optimization for Lightweight Medical AI

AutoML frameworks can automate:

- Hyperparameter tuning
- Model pruning
- Quantization

Leading to even more optimized, deployable models.

VII. CONCLUSION

Alzheimer's Disease poses one of the most pressing global healthcare challenges due to its progressive nature and the difficulty of detecting early cognitive decline. Traditional diagnostic procedures rely heavily on radiological expertise and manual MRI interpretation, which are subjective, time-consuming, and infeasible at population scale. Deep learning has shown immense potential to overcome these limitations, but most existing models depend on computationally intensive architectures that hinder real-world deployment.

This research presents a **lightweight, interpretable, and clinically practical deep-learning framework** based on MobileNetV2, enhanced with Grad-CAM visualization. Through transfer learning, optimized preprocessing, and systematic fine-tuning, the model successfully classifies four dementia stages with strong accuracy while maintaining a minimal computational footprint. More importantly, the Grad-CAM module highlights anatomically plausible brain regions associated with AD pathology, thereby bridging the gap between automated predictions and clinical reasoning.

The model's small size, fast inference speed, and CPU-friendly design make it suitable for deployment in resource-limited settings—rural hospitals, mobile health units, telemedicine platforms, and edge-AI healthcare devices. The integration of interpretability ensures higher clinician trust, making the approach not only technically sound but also medically relevant. In summary, the proposed system represents a meaningful step toward **scalable, transparent, and affordable AI-based Alzheimer's screening** and lays a foundation for future clinical applications.

ACKNOWLEDGMENT

It gives us great pleasure in presenting the preliminary project report on 'Lightweight Deep Learning-Based Alzheimer's Detection Using MobileNetV2 and Grad-CAM: A Comprehensive Review'. I would like to take this opportunity to thank our internal guide and Assistant Professor, Poonam Singh for giving us all the help and guidance we needed. We are really grateful to him for his kind support. Her valuable suggestions were very helpful.

REFERENCES

[1]. A. Alafif, M. Abuhaija, and S. Masoud, "Deep learning approaches for Alzheimer's disease classification using MRI: A systematic review," Diagnostics, vol. 14, no. 3, pp. 1–22, 2024.

[2]. A. Assaduzzaman, M. R. H. Mondal, and M. M. Rahman, "Lightweight convolutional architectures for Alzheimer's disease detection using structural MRI," Informatics in Medicine Unlocked, vol. 40, p. 101115, 2024.

[3]. Bhattacharya, D. Chaudhuri, and R. Sanyal, "Explainable AI with Grad-CAM for neurodegenerative disease diagnosis: A comprehensive survey," Computers in Biology and Medicine, vol. 178, p. 106574, 2024.

[4]. Fan, M. Batmanghelich, C. Davatzikos, and D. Shen, "Multimodal classification of Alzheimer's disease combining MRI and PET using relevance vector machines," NeuroImage, vol. 41, no. 2, pp. 559–568, 2008.

[5]. Jin, K. Xue, and L. Chen, "Transfer learning for early Alzheimer's detection using MobileNet variants," Journal of Imaging, vol. 10, pp. 1–12, 2024.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, February 2026
Impact Factor: 8.2

[6]. Shi, H. Wang, and Y. Zhang, "Deep convolutional learning for AD vs. MCI classification from 2D and 3D MRI representations," PLOS One, vol. 19, no. 4, pp. 1–18, 2024.
[7]. Hammad, R. Abbas, A. Salem, and A. El-Feshawy, "Explainable Alzheimer's disease classification using Grad-CAM with efficient deep networks," Diagnostics, vol. 13, no. 4, pp. 1–15, 2023.
[8]. D. Shen, G. Wu, and H.-I. Suk, "Deep learning in medical image analysis," Annual Review of Biomedical Engineering, vol. 19, pp. 221–248, 2017.
[9]. Chen, W. Li, and J. Zhao, "Hybrid CNN-Transformer networks for Alzheimer's disease staging," IEEE Access, vol. 12, pp. 118503–118520, 2024.
[10]. Litjens et al., "A survey on deep learning in medical imaging," Medical Image Analysis, vol. 42, pp. 60–88, 2017.
[11]. Rajagopal, M. Pandey, and S. R. Patel, "MobileNetV2-based AD–MCI–HC classification using lightweight MRI models," Scientific Reports, vol. 15, pp. 1–15, 2025.
[12]. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.
[13]. Shalaby, M. Abou El-Magd, and A. Soliman, "Deep feature fusion for Alzheimer's diagnosis using 3D MRI," Diagnostics, vol. 14, no. 2, pp. 1–15, 2024.
[14]. Brownlee, "A gentle introduction to transfer learning," Machine Learning Mastery, 2021.
[15]. J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in Proc. IEEE CVPR, 2018, pp. 7132–7141.
[16]. J. Sheikh, T. Kawther, and R. Ahmed, "Benchmarking deep architectures for Alzheimer's disease progression prediction," Diagnostics, vol. 15, pp. 1–16, 2025.
[17]. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE CVPR, 2016, pp. 770–778.
[18]. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv:1409.1556, 2014.
[19]. Abrol et al., "Deep neural networks for Alzheimer's detection using 3D MRI," Human Brain Mapping, vol. 44, no. 1, pp. 45–63, 2023.
[20]. Bäckström, A. Berg, and L. Nyström, "Explainable AD diagnosis using attention-guided CNNs," Scientific Reports, pp. 1–12, 2024.
[21]. M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for CNNs," in Proc. ICML, 2019, pp. 6105–6114.
[22]. M. Zhou, X. Li, and Z. Chen, "Unified deep multimodal framework for Alzheimer's diagnosis with explainability," Scientific Reports, vol. 15, pp. 1–12, 2025.
[23]. R. Klöppel et al., "Automatic classification of MR scans in Alzheimer's disease using SVM," Brain, vol. 131, no. 3, pp. 681–689, 2008.
[24]. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, pp. 1735–1780, 1997.
[25]. S. Ioffe and C. Szegedy, "Batch normalization," in Proc. ICML, 2015, pp. 448–456.
[26]. S. Shanmugam, A. Kumar, and P. D. Singh, "Vision Transformers for neurodegenerative disease screening," IEEE Access, vol. 11, pp. 145443–145460, 2023.
[27]. S. Woo, J. Park, J. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in Proc. ECCV, 2018, pp. 3–19.
[28]. T. Hussain et al., "Deep-learning-based Alzheimer's detection using structural MRI with extensive benchmarking," Scientific Reports, vol. 15, pp. 1–16, 2025.
[29]. T. Xiao, Y. Wang, and L. Qiang, "Cross-attention networks for early AD detection," IEEE Journal of Biomedical and Health Informatics, vol. 28, pp. 1–12, 2024.
[30]. V. Gulshan et al., "Interpretable computer vision for medical diagnosis," Nature Medicine, pp. 1–12, 2020.
[31]. V. Srinivasan et al., "Mobile health and edge-AI solutions for neurodegenerative diseases," IEEE Internet of Things Journal, vol. 10, pp. 1–15, 2023.