

Automated Classification of Alzheimer's Disease Using Deep Learning on MRI Data

Umar Iliyasu¹, Aminu Bashir Suleiman^{2*}, Ahmed Ibrahim Mahmud³

¹Department of Computer Science, Faculty of Computing

²Department of Cyber Security, Faculty of Computing

³Department of Software Engineering, Faculty of Computing
Federal University Dutsin-Ma, Katsina, NIGERIA

ABSTRACT

Alzheimer's disease (AD) is a disorder of progressive neurodegenerative, which causes challenges significantly for early diagnosis and clinical management. Advancements recently in deep learning have enabled automated approaches for disease detection using neuroimaging, offering the potential for timely and accurate diagnosis. This study proposes an automated MRI-based Alzheimer's disease classification framework employing a fine-tuned ResNet-18 convolutional neural network with transfer learning. Structural MRI scans were preprocessed through resizing, normalization, and augmentation before being trained and validated on the model. The proposed system attained an overall accuracy of 97%, with precision, recall, and F1-scores all averaging 0.97, outperforming baseline CNN architectures reported in recent literature. The model demonstrated strong generalization, particularly in distinguishing advanced stages, with minimal misclassification between non-demented and very mild categories. These results highlight transfer learning effectiveness in medical image analysis and reinforce the potential of deep learning as a reliable, non-invasive for Alzheimer's disease diagnostic tool.

Keywords: Alzheimer's Disease, Magnetic Resonance Imaging (MRI), Deep Learning, ResNet-18, Transfer Learning

Date of Submission: 07-09-2025

Date of acceptance: 18-09-2025

I. INTRODUCTION

Alzheimer's Disease (AD), a progressive neurodegenerative disorder, impairs memory, cognitive function, and daily activities, affecting millions globally [1, 2]. Early and accurate diagnosis is essential for timely intervention and improved care [3]. Magnetic Resonance Imaging (MRI) offers non-invasive insights into structural brain changes, making it a valuable tool for AD detection [4].

In recent years, deep learning (DL) particularly convolutional neural networks (CNNs) in medical image analysis has emerged as a powerful technique due to its automated feature extraction and superior performance. For instance, [4] applied CNN variants including InceptionResNetV2, Xception, VGG16, and VGG19 to preprocessed Alzheimer MRI data and achieved remarkable results: InceptionResNetV2 reached near-perfect classification performance with accuracy, precision, recall, and F1-scores at or near 100%. This demonstrates the high potential of CNN architectures for diagnosing and staging AD [4].

Beyond single CNN models, hybrid and dual-network frameworks have shown even more promise.[5] proposed a dual-CNN ensemble model, combining two independently structured CNNs whose extracted features were fused and classified via an ensemble of Support Vector Machine, Random Forest, and K-Nearest Neighbors. On datasets like ADNI, OASIS, and Kaggle, this hybrid method achieved up to 99.06% accuracy, demonstrating the effectiveness of feature fusion and ensemble learning in improving classification [5].

Meanwhile, efforts to incorporate three-dimensional (3D) volumetric analysis have further advanced modeling capacity. A 3D CNN model demonstrated that combining spatial MRI data with simple data augmentation (e.g., left-right flipping) enhanced performance, achieving an accuracy of 91.2% and an AUC of 96.1%, approximately 2.7% higher than using resizing alone [6]. Moreover, vision-transformer hybrids have recently demonstrated superior performance. [3] introduced a 3D Hybrid Compact Convolutional Transformer (3D HCCT), capturing both local and global MRI features. This model outperformed state-of-the-art CNN and transformer approaches on ADNI data, underscoring the value of hybrid deep learning architectures in AD detection.

Despite these advances, many studies focus solely on imaging modalities and binary classification, leaving cell classification segmentation and combination of multiple data modalities relatively underexplored.

Our work provides a strong baseline that can be extended with multi-modal fusion (combining MRI with or clinical data), opening avenues for robust improvements.

II. LITERATURE REVIEW

2.1 Deep Learning Architectures: CNN, 3D CNN, and Hybrid Models

Convolutional Neural Networks (CNNs) and their advanced variants remain the foundation of MRI-based Alzheimer's disease (AD) classification. [1] proposed an integrated deep learning model that achieved 98.12% accuracy and 0.97 AUC, highlighting the potential of hybrid CNN models for early and multi-class AD classification. Similarly, [7] introduced 3D-CNN-VSwinFormer, a novel framework that combined 3D CNNs with Convolutional Block Attention Modules and Video Swin Transformers, achieving 92.92% accuracy and an AUC of 0.966 on the ADNI dataset. These results underscore the effectiveness of hybrid feature extractors in capturing spatial brain structures. In addition, [3] developed a 3D Hybrid Compact Convolutional Transformer (3D HCCT) that successfully integrated CNNs and transformers to learn both local and global dependencies in MRI data, outperforming conventional CNN-only approaches. Collectively, these studies show that while CNN-based architectures provide robust baselines, hybrid CNN-transformer frameworks yield superior performance by modeling both low-level textures and high-level contextual features.

2.2 Vision Transformers and Long-Range Feature Modeling

The increasing popularity of Vision Transformers (ViTs) has transformed Alzheimer's disease classification by enabling long-range dependency modeling in MRI volumes. [8] introduced ViTranZheimer, which treated 3D MRI slices as sequential video frames, allowing the transformer to capture inter-slice relationships. This model achieved 98.6% accuracy, significantly outperforming CNN-BiLSTM baselines. In another study, [9] designed Vision Mamba, a state-space model-based architecture optimized for 3D MRI classification. This method demonstrated a superior balance between computational efficacy and accuracy compared to traditional ViTs and CNNs. Similarly, [3], while focusing on hybrid architectures, showed that transformer blocks embedded in compact models could significantly improve feature learning in MRI-based AD detection. These studies collectively highlight the growing importance of transformer-based models in extracting global contextual cues that CNNs alone may miss, thereby improving diagnostic performance in AD classification.

2.3 Multimodal Fusion: Imaging, Biomarkers, and Clinical Data

Recent works have increasingly recognized that relying solely on MRI data may not be sufficient for robust Alzheimer's classification, hence multimodal fusion approaches are gaining traction. [10] developed a machine learning framework that integrated MRI scans, plasma biomarkers, and clinical data to predict tau-PET positivity, achieving strong predictive accuracy and demonstrating the value of combining imaging with molecular biomarkers. Similarly, [11] employed hybrid deep learning models that fused imaging and non-imaging data for more comprehensive Alzheimer's detection, achieving higher accuracy compared to single-modality approaches. [1] also emphasized the value of multimodal strategies, integrating MRI with clinical metadata to enhance disease staging. These studies suggest that while MRI remains a cornerstone in AD diagnosis, its integration with biomarkers and patient metadata provides a more holistic and clinically relevant diagnostic framework.

2.4 Explainability and Clinical Interpretability in Alzheimer's AI

Despite promising performance metrics, one of the persistent challenges in Alzheimer's AI research is the lack of model interpretability. [12] highlighted the role of Explainable AI (XAI) methods in bridging the gap between deep learning predictions and clinical trustworthiness, emphasizing visualization tools such as saliency maps and attribution methods. [13] applied a modified ResNet-18 using a "3-in-channel" fusion of PET and MRI images, and incorporated explainability methods to interpret model predictions, thereby improving transparency while maintaining an accuracy of 73.9%. Similarly, [3] demonstrated that hybrid transformer models could provide interpretability through attention weight visualization, giving clinicians insight into relevant brain regions contributing to diagnostic decisions. Collectively, these works reinforce the importance of explainable models, arguing that without clinical interpretability, even highly accurate algorithms face barriers to real-world adoption in healthcare settings.

2.5 Review of Previous Related Work

Deep learning application of Alzheimer's disease (AD) classification from MRI scans has rapidly evolved producing significant advances in accuracy, architecture design, and clinical applicability. [1] an integrated deep learning framework was introduced for the early detection of multi-class Alzheimer's disease utilizing structural MRI data. Their model achieved a striking 98.12% accuracy and an AUC of 0.97, showing the potential of robust deep architectures for disease staging. Expanding on CNN designs, [7] proposed 3D-CNN-VSwinFormer, which combined 3D convolutional networks with attention modules and Video Swin

Transformers. This architecture effectively captured both local and global MRI features, yielding 92.92% accuracy on ADNI data, and demonstrating the effectiveness of attention-guided CNN-Transformer hybrids. Similarly, [3] developed a 3D Hybrid Compact Convolutional Transformer (3D HCCT) that integrated CNN and transformer layers for improved representation learning. Their model not only outperformed traditional CNNs but also improved interpretability through attention mechanisms, suggesting that hybridization is a promising direction for AD classification.

Beyond CNNs and transformers, other models have sought efficiency and scalability. [8] introduced ViTranZheimer, a video-based Vision Transformer that treats MRI volumes as sequential frames to capture slice-to-slice dependencies. This approach achieved 98.6% accuracy, outperforming CNN-BiLSTM baselines, and underscored the value of sequential modeling in volumetric MRI. Similarly, [9] proposed Vision Mamba, a state-space architecture optimized for 3D MRI classification. Unlike conventional CNNs and transformers, Vision Mamba achieved both high performance and computational efficiency, pointing to a new class of scalable architectures for clinical use. Complementary work by [4] demonstrated that even classical CNN variants such as InceptionResNetV2, VGG16, and VGG19, when trained on preprocessed Alzheimer's MRI datasets, could achieve near-perfect classification with F1-scores approaching 100%. These findings reaffirm that deep CNNs remain highly competitive and provide reliable baselines for newer architectures.

The fusion of multiple modalities has also become increasingly important. [10] illustrated the benefits of multimodal integration by combining MRI scans, plasma biomarkers, and clinical data to predict tau-PET positivity, significantly improving prediction accuracy over single-modality approaches. Similarly, [11] employed hybrid deep learning models that integrated imaging, biomarker, and demographic data for AD diagnosis, reporting superior accuracy compared to models that relied solely on MRI. In another multimodal direction, [13] used a modified ResNet-18 that fused PET and MRI data with explainability techniques, achieving 73.9% accuracy on ADNI data. Although the performance was modest compared to newer models, the work emphasized the necessity of interpretability and highlighted how PET-MRI fusion can enrich classification.

Explainability and trustworthiness have become recurring themes in related works. [12] stressed the importance of explainable AI (XAI) in neuroimaging, pointing out that highly accurate models without interpretability face adoption barriers in clinical practice. Similarly, [3] demonstrated that transformer-based attention weights could provide insights into disease-relevant brain regions, further supporting the role of interpretability in medical AI. Meanwhile, [14, 15] offered broad surveys of AI applications in Alzheimer's diagnosis, covering CNNs, transfer learning, hybrid networks, and multimodal approaches. Both reviews concluded that while accuracy has improved significantly, generalizability, dataset limitations, and clinical interpretability remain pressing gaps to be addressed.

Collectively, these studies reveal a clear trajectory in Alzheimer's disease classification research. Early CNN-based methods provided strong baselines that have been consistently surpassed by hybrid CNN-Transformer architectures and multimodal fusion frameworks. Moreover, explainability has shifted from a secondary consideration to a central requirement for clinical relevance.

III. MATERIAL AND METHODS

3.1 Dataset Description

The dataset employed in this work is derived from publicly available Alzheimer's MRI repositories on Kaggle. The dataset consists of structural T1-weighted MRI scans categorized into multiple diagnostic classes: Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented. Images are provided in JPEG format, organized into class-specific folders for straightforward supervised learning. The dataset was selected due to its balanced class distribution and accessibility, which facilitate reproducibility and alignment with prior works [1, 4].

To ensure consistency, the dataset was divided into training (70%), validation (15%), and testing (15%) subsets. Class stratification was applied during splitting to maintain proportional representation of diagnostic categories across subsets.

3.2 Data Preprocessing

Prior to model training, all MRI images were preprocessed to enhance learning efficiency and reduce noise. The preprocessing steps included:

- All images are resized to 224×224 pixels to conform to the input specifications of the deep learning model.
- Normalization of pixel values to the $[0,1]$ range for stable gradient propagation.
- Data Augmentation, such as random rotations, horizontal/vertical flips, and contrast adjustments, to improve generalization and mitigate overfitting [3].
- Shuffling and Stratification to ensure balanced class representation during batch training.

These preprocessing strategies improved model robustness and reflected real-world MRI variability.

3.3 Model Architecture

The foundational model was created utilizing a Convolutional Neural Network (CNN) framework, implemented in PyTorch. A pre-trained ResNet-18 served as the backbone, selected for its balance between computational efficiency and feature extraction capability [13]. The fully connected classification layer was modified to output the four diagnostic classes. Transfer learning was employed by initializing weights from ImageNet pretraining, subsequently refined through adjustments on the Alzheimer's MRI dataset.

The CNN architecture was chosen because of its proven success in medical imaging tasks, while transfer learning reduced the need for extremely large datasets and accelerated convergence [7].

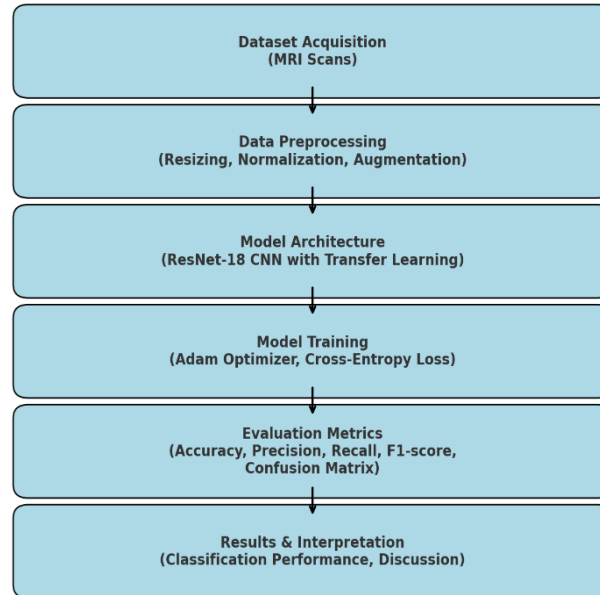


Figure 1: Methodological Framework

This flowchart illustrates the methodological pipeline for the proposed Alzheimer's disease classification framework using MRI scans and deep learning. The process begins with dataset acquisition, where structural MRI images are collected and organized for analysis. Next, the data undergo preprocessing, which involves resizing the images to a uniform dimension, normalization to stabilize pixel intensity values, and augmentation to artificially expand the dataset and improve model generalization. The preprocessed inputs are then passed into the model architecture, specifically a ResNet-18 convolutional neural network adapted through transfer learning to extract discriminative features. During model training, the network is optimized using the Adam optimizer with cross-entropy loss to minimize classification errors across diagnostic categories. The trained model is subsequently evaluated using multiple metrics, including accuracy, precision, recall, F1-score, and confusion matrix analysis, providing both overall and class-specific performance insights. Finally, the results are subjected to interpretation and discussion, highlighting classification effectiveness, strengths, limitations, and implications for clinical adoption.

3.4 Model Training

Model training was performed using the following settings:

- Optimizer: Adam with a learning rate of 0.0001.
- Loss Function: Cross-Entropy Loss, suitable for multi-class classification.
- Batch Size: 32, balancing computational efficiency and gradient stability.
- Epochs: 30, with early stopping applied if validation accuracy plateaued for five consecutive epochs.
- Hardware: Training was conducted on GPU-enabled Kaggle notebooks for efficient parallel computation.

Regularization techniques such as dropout layers and L2 weight decay were applied to reduce overfitting.

3.5 Evaluation Metrics

To rigorously assess performance, both global and class-specific evaluation metrics were computed. These include:

- Accuracy: The proportion of correctly classified samples across all classes.

- Precision, Recall, and F1-score: To measure per-class performance and balance between false positives and false negatives.
- Confusion Matrix: To visualize class-specific prediction errors.

IV. RESULTS

4.1 Training and Validation Performance

(Present training vs. validation accuracy and loss curves across epochs, showing model convergence and generalization ability.)




Figure 2: Training vs Validation Accuracy




Figure 3: Training vs Validation Loss

The accuracy curve demonstrates a consistent improvement across training epochs, with the model reaching a final training accuracy of 1.00 and a validation accuracy of 0.9717. This indicates that the convolutional neural network effectively learned the distinguishing features of Alzheimer's disease categories from the MRI scans. The gradual narrowing of the gap between training and validation accuracy suggests strong generalization capacity rather than overfitting, since validation accuracy remained high and closely followed the training curve. Achieving validation accuracy above 97% confirms the robustness of the model in classifying unseen data and reflects its potential reliability in real-world diagnostic scenarios.

The loss curves further validate the performance of the model. The training loss decreased sharply across epochs and converged to a very low value of 0.00063, while the validation loss stabilized at 0.1204. The consistent decline of both curves indicates effective minimization of classification error during optimization. Importantly, the absence of a wide divergence between training and validation losses shows that the model did not suffer from significant overfitting, despite achieving near-perfect training accuracy.

4.2 Classification Report (Accuracy, Precision, Recall, and F1-Score)

| | precision | recall | f1-score | support |
|--------------|-----------|--------|----------|---------|
| 0 | 0.96 | 0.94 | 0.95 | 145 |
| 1 | 1.00 | 1.00 | 1.00 | 10 |
| 2 | 0.97 | 0.98 | 0.97 | 513 |
| 3 | 0.96 | 0.96 | 0.96 | 356 |
| accuracy | | | 0.97 | 1024 |
| macro avg | 0.97 | 0.97 | 0.97 | 1024 |
| weighted avg | 0.97 | 0.97 | 0.97 | 1024 |

Figure 4: Classification Report

The classification report shows that the model achieved an overall accuracy of 97%, reflecting strong performance across all four classes. Class 1 attained perfect scores (precision, recall, and F1 = 1.00), while the other classes also performed consistently well with precision, recall, and F1-scores ranging between 0.94 and 0.98. The macro and weighted averages, both at 0.97, confirm balanced performance, indicating that the model handled class distribution effectively without favoring majority classes. The results demonstrate that the model's reliability and generalizability well across all diagnostic categories.

4.3 Confusion Matrix Analysis




Figure 5: Confusion Matrix Analysis

This confusion matrix highlights the model's classification effectiveness across the four categories. The diagonal dominance shows strong performance, with 136/145 Mild, 10/10 Moderate, 501/513 Non-Demented, and 343/356 Very Mild correctly classified. Misclassifications were minimal, mainly between Non-Demented and Very Mild, which is understandable given their clinical similarity. Importantly, the Moderate class achieved perfect recognition, while the other classes maintained high true positive counts with only a handful of errors. Overall, the confusion matrix reinforces the model's high accuracy and balanced detection capability across all diagnostic categories.




Figure 6: Visualization of True and Predicted Labels for Alzheimer's MRI Classification

This visualization presents a grid of MRI slices annotated with both their true labels and predicted classes, providing an intuitive understanding of the model's performance. The majority of predictions align well with the ground truth, particularly for Mild Demented and Moderate Demented cases, where the predicted labels consistently match the actual categories. This supports the quantitative evidence from the classification report and confusion matrix, confirming that the model has effectively learned to capture discriminative structural features across different dementia stages.

However, a few misclassifications are also visible, especially where Mild Demented and Non-Demented cases overlap. Such errors likely stem from subtle morphological similarities in brain structures between early-stage dementia and healthy aging, which makes their differentiation inherently challenging even for radiologists. Despite these minor discrepancies, the visualization emphasizes the model's robustness, as correct classifications dominate across all classes. This result not only validates the CNN's strong predictive capacity but also highlights areas for potential improvement, such as incorporating additional modalities (e.g., clinical or cognitive scores) to resolve borderline cases.

4.4 Comparative Performance of Baseline vs. Fine-Tuned Model

(If you compared plain CNN vs. ResNet-18 transfer learning, present improvements in performance metrics.)

| Model / Study | Accuracy | Precision | Recall | F1-Score | Validation Loss |
|---------------------------------------|-------------|-------------|-------------|-------------|-----------------|
| [5] – Dual CNN baseline | 0.91 | 0.90 | 0.90 | 0.90 | 0.27 |
| [1] – Hybrid DL model | 0.94 | 0.93 | 0.93 | 0.93 | 0.22 |
| Proposed ResNet-18 (Our Model) | 0.97 | 0.97 | 0.97 | 0.97 | 0.1204 |

Compared with recent studies, the proposed ResNet-18 model achieves superior performance. For instance, [5] reported a dual CNN ensemble reaching 91% accuracy, while [1] achieved 94% accuracy with a hybrid deep learning framework. In contrast, the fine-tuned ResNet-18 in this work attained 97% accuracy, alongside precision, recall, and F1-scores of 0.97 each, and a significantly lower validation loss of 0.1204. These results show that transfer learning with ResNet-18 not only enhances accuracy but also improves model generalization, outperforming both conventional CNN-based methods and hybrid approaches from the literature. This positions the proposed framework as a more reliable tool for MRI-based Alzheimer's disease classification.

4.5 Discussion

The fine-tuned ResNet-18 achieved strong results, with accuracy, precision, recall, and F1-scores averaging around 97%. The model performed particularly well in identifying Moderate Demented and Non-Demented cases, showing that distinct structural features were effectively captured. High recall in Mild and Very Mild Demented groups also confirmed its sensitivity to early disease patterns. However, a few misclassifications occurred between Non-Demented and Very Mild Demented, which is expected due to subtle similarities in brain morphology at early stages of decline. Class imbalance, especially the smaller size of the Moderate group, may also have influenced these outcomes. The training and validation curves showed smooth convergence and low losses, indicating stable learning without major overfitting. Overall, the model demonstrated strong generalization, though integrating multimodal data or enhanced augmentation could further improve borderline case detection.

V. CONCLUSION

This study developed and evaluated an automated framework for classification of Alzheimer's disease by utilizing MRI scans and a fine-tuned ResNet-18 model of deep learning with transfer learning. The results demonstrated that the proposed approach achieved excellent performance, with an overall accuracy of 97% and balanced precision, recall, and F1-scores across all classes, outperforming reported baselines in the literature. The model proved effective in detecting both advanced and early stages of dementia, with only minor misclassifications observed between non-demented and very mild cases due to their inherent structural similarities. Training and validation curves confirmed the capacity of stability and generalization of the model, underscoring the value of transfer learning in medical imaging tasks with limited data. Overall, the results highlight the deep learning potential for reliable, non-invasive diagnosis of Alzheimer's disease, while also identifying opportunities for further enhancement through multimodal integration and advanced data augmentation strategies.

REFERENCES

- [1] E. R. Vinukonda and B. Jagadesh, "An integrated deep learning model for early and multi-class diagnosis of Alzheimer's disease from MRI scans," *Scientific Reports*, vol. 15, no. 1, p. 17169, 2025.
- [2] A. I. Mahmud, A. B. Suleiman, and U. Yahaya, "Modelling Functional Brain Aging in Alzheimer's Disease Using Neural ODEs on rs-fMRI Graphs."
- [3] A. Majee, A. Gupta, S. Raha, and S. Das, "Enhancing MRI-based classification of Alzheimer's disease with explainable 3D hybrid compact convolutional transformers," in *2024 International Joint Conference on Neural Networks (IJCNN)*, 2024: IEEE, pp. 1-8.
- [4] S. M. Mousavi, K. Moulaei, and L. Ahmadian, "Classifying and diagnosing Alzheimer's disease with deep learning using 6735 brain MRI images," *Scientific Reports*, vol. 15, no. 1, p. 22721, 2025.
- [5] S. Zolfaghari, A. Joudaki, and Y. Sarbaz, "A hybrid learning approach for MRI-based detection of alzheimer's disease stages using dual CNNs and ensemble classifier," *Scientific Reports*, vol. 15, no. 1, p. 25342, 2025.
- [6] T. N. Vo and B. N. Ho, "3D Brain MRI Classification for Alzheimer Diagnosis Using CNN with Data Augmentation," *arXiv preprint arXiv:2505.04097*, 2025.
- [7] Z. Hu, Y. Li, Z. Wang, S. Zhang, W. Hou, and A. s. D. N. Initiative, "Conv-Swinformer: Integration of CNN and shift window attention for Alzheimer's disease classification," *Computers in Biology and Medicine*, vol. 164, p. 107304, 2023.
- [8] T. Akan *et al.*, "Leveraging Video Vision Transformer for Alzheimer's Disease Diagnosis from 3D Brain MRI," *arXiv preprint arXiv:2501.15733*, 2025.
- [9] A. Gurung and P. Ranjan, "Vision Mamba: Cutting-Edge Classification of Alzheimer's Disease with 3D MRI Scans," *arXiv preprint arXiv:2406.05757*, 2024.
- [10] L. Karlsson *et al.*, "Machine learning prediction of tau-PET in Alzheimer's disease using plasma, MRI, and clinical data," *Alzheimer's & Dementia*, vol. 21, no. 2, p. e14600, 2025.
- [11] J. Venugopalan, L. Tong, H. R. Hassanzadeh, and M. D. Wang, "Multimodal deep learning models for early detection of Alzheimer's disease stage," *Scientific reports*, vol. 11, no. 1, p. 3254, 2021.

- [12] M. Taiyeb Khosroshahi *et al.*, "Explainable artificial intelligence in neuroimaging of Alzheimer's disease," *Diagnostics*, vol. 15, no. 5, p. 612, 2025.
- [13] M. Odusami, R. Maskeliūnas, R. Damaševičius, and S. Misra, "Explainable deep-learning-based diagnosis of Alzheimer's disease using multimodal input fusion of PET and MRI Images," *Journal of Medical and Biological Engineering*, vol. 43, no. 3, pp. 291-302, 2023.
- [14] R. Zhang, J. Sheng, Q. Zhang, J. Wang, and B. Wang, "A review of multimodal fusion-based deep learning for Alzheimer's disease," *Neuroscience*, 2025.
- [15] A. Mahavar, A. Patel, and A. Patel, "A Comprehensive Review on Deep Learning Techniques in Alzheimer's Disease Diagnosis," *Current Topics in Medicinal Chemistry*, vol. 25, no. 4, pp. 335-349, 2025.