International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

NeuroStageAI - AI Powered Alzheimer's Disease Detection

Farheen Sultana[1], Mohammed Hannan Ahmed[2], Mohd Abdul Rafay Kamran[3], Syed Zohair Sayeed[4]*

[1] *Assistant Professor,[B.Tech,M.Tech(PhD)]*

[2] *Department of IT, Nawab Shah Alam Khan College of Engineering and Technology, Hyderabad, India*

ABSTRACT

Alzheimer's Disease (AD) is a severe neurodegenerative condition that leads to a gradual decline in memory, thinking, behavior, and daily living skills. Its slow onset and gradual worsening make it crucial to diagnose early and classify the stage correctly for effective treatment, patient care planning, and therapeutic choices. This study uses recent progress in brain imaging, artificial intelligence, and online computing to present a detailed system for automatically detecting Alzheimer's at various stages and predicting its progression.

The system is based on structural MRI brain scans from public datasets, which are divided into training (70%), validation (15%), and test (15%) groups. It includes different categories: "Very Mild Demented," "Mild Demented," "Moderate Demented," and "Non Demented." To reduce overfitting and improve model performance in medical imaging, we applied various data enhancement methods like random rotations, horizontal and vertical flips, and adjustments to contrast and brightness.

Two deep-learning models were fine-tuned on the MRI data: MobileNet, a lightweight transfer-learning model, and InceptionV3, a deeper network. MobileNet reached a training accuracy of nearly 99% and a test accuracy of about 98.35%. InceptionV3 achieved roughly 99.48% training accuracy and around 94.25% test accuracy. The web interface, created with Flask (using Python for the backend and HTML/CSS/JavaScript for the front end), allows clinicians or users to upload MRI scans and receive predictions about the AD stage while visualizing model-explanation overlays.

An innovative module uses intermediate feature embeddings from MobileNet to estimate the chance of disease progression over the next 12-24 months (stable versus deteriorating), which offers a prognostic aspect beyond simple classification. We enhance explainability using Grad-CAM heat maps to show which brain areas influence the model's decisions, promoting trust and transparency in clinical settings.

This system can operate on basic hardware (like an Intel Pentium i3 with 8 GB RAM and 20 GB storage), making it suitable for educational institutions, small clinics, or pilot programs. The combined framework for classification and progression prediction displays high accuracy, good usability, and clear outputs. This marks a significant advancement in AI-assisted clinical support for Alzheimer's in situations with limited resources. The added progression-forecasting module increases the system's usefulness from static diagnosis to dynamic prognosis, encouraging timely interventions and personalized patient care.

Keywords: Alzheimer's Disease, Deep Learning, MRI, MobileNet InceptionV3, Data Augmentation, Web Interface

Introduction

Alzheimer's Disease (AD) stands among the most pressing public-health challenges of our time. As a progressive neurodegenerative disorder primarily affecting older adults, AD leads to gradual yet irreversible declines in memory, cognition, behaviour, and the capacity for independent daily living. Many individuals transition from a stage of Mild Cognitive Impairment (MCI) to full-blown dementia, often experiencing subtle symptoms long before clinical diagnosis. Given the increasing global elderly population, timely detection and

accurate stagewise classification of AD have become critical to enable early interventions, slow disease progression, support care planning and improve patient outcomes.

Neuroimaging has emerged as a cornerstone in the clinical and research evaluation of Alzheimer's disease. In particular, structural brain imaging modalities—most notably magnetic resonance imaging (MRI)—offer a non-invasive window into brain morphology and pathology. MRI studies have repeatedly shown that hallmark features of AD include atrophy of medial temporal lobe structures (such as the hippocampus and entorhinal cortex), ventricular enlargement, cortical thinning and white matter hyperintensities. In addition to structural changes, research has shown that MRI can assist in excluding other causes of cognitive decline (such as vascular dementia or hydrocephalus) and thus plays a vital diagnostic role.

Nonetheless, traditional MRI interpretation relies heavily on expert radiological assessment and remains less suitable for large-scale, automated screening.

In parallel, the rapid evolution of artificial intelligence (AI), and in particular deep learning (DL) techniques, has transformed the landscape of medical image analysis. Convolutional neural networks (CNNs), autoencoders, recurrent networks and hybrid architectures have shown substantial promise in automatically extracting complex features from high-dimensional neuroimaging data. A systematic review of DL applications in AD classification found accuracies of up to ~96% using MRI only, and higher when combining modalities such as PET and fluid biomarkers. Moreover, recent studies suggest that deep learning on MRI may detect subtle neurodegenerative changes well before overt clinical symptoms manifest—opening the possibility of earlier diagnosis and intervention. However, despite these advancements, major challenges remain: class imbalance in datasets, limited availability of large labelled cohorts, model interpretability ("black box" concerns), and the gap between high-performance research prototypes and deployable clinical systems.

To address these challenges, our project presents a practical, user-accessible web-based system for multi-stage Alzheimer's detection from MRI scans, integrated with a novel prognostic component. Unlike many prior works that focus on binary classification (e.g., AD vs non-AD) or on offline analysis, our system supports four AD stages—Non Demented, Very Mild Demented, Mild Demented, Moderate Demented—and is built for deployment in resource-constrained environments (e.g., modest hardware). We employ efficient architectures (such as MobileNet and InceptionV3) coupled with data-augmentation to tackle overfitting and class imbalance, and integrate explainability (via Grad-CAM) to increase clinician trust. Crucially, we extend beyond static classification by adding a progression prediction module: using extracted feature embeddings, the system estimates the risk of disease progression over the next 12-24 months, thereby providing clinicians with actionable prognostic insight rather than only a snapshot diagnosis.

In summary, this study aims to bridge the gap between advanced neuroimaging-AI research and practical clinical/educational deployment. By combining efficient deep-learning models, deployable web architecture, multi-stage classification and prognostic forecasting, we target an accessible, interpretable and clinically relevant solution for early Alzheimer's detection and monitoring. The following sections describe the system analysis and design, methodology, results, conclusion and future enhancements of our work.

Nomenclature

AD	**Alzheimer's Disease**
AI	**Artificial Intelligence**
CNN	**Convolutional Neural Network**
DL	**Deep Learning**
MRI	**Magnetic Resonance Imaging**
MCI	**Mild Cognitive Impairment**
RAM	**Random Access Memory**
API	**Application Programming Interface**
CPU	**Central Processing Unit**
IP	**Internet Protocol**

2. Literature Review

Alzheimer's Disease (AD) has been widely studied using neuroimaging, especially structural MRI, because it captures brain atrophy patterns associated with disease progression. Several systematic reviews highlight that machine learning and deep learning methods on MRI and PET data have significantly improved early AD detection, but also point out challenges such as limited dataset size, class imbalance and lack of clinical deployment. PMC+2MDPI+2

Early works mainly relied on conventional machine learning pipelines, where handcrafted features (volumetric measures, cortical thickness, texture descriptors) were extracted from MRI images and then fed to classifiers like SVM, Random Forest or k-NN. Over time, deep learning—especially convolutional neural networks (CNNs)—became dominant due to their ability to automatically learn hierarchical features from raw or minimally processed MRI volumes. Studies using 3D CNNs and variants such as 3D AlexNet and 3D ResNet have reported substantial gains in accuracy for distinguishing normal controls, Mild Cognitive Impairment (MCI), and AD, with some models achieving competitive performance across multiple public datasets.

Recent research has increasingly focused on **multi-class classification** rather than simple binary AD vs. non-AD problems. For example, some works employ deep CNNs like InceptionV3 to classify subjects into four risk or severity levels based on MRI, including Non-Demented, Very Mild Demented, Mild Demented and Moderate Demented categories, achieving promising but dataset-dependent performance. Health Informatics Journal Other frameworks extract deep features from networks such as EfficientNet and MobileNet, then fuse them using techniques like canonical correlation analysis to improve multi-stage dementia classification accuracy. However, most of these systems primarily focus on **static diagnosis**; they rarely model future progression (e.g., stable vs. worsening over time) and often assume access to high-end GPU resources.

Alongside CNNs, transformer-based architectures have recently gained attention in AD imaging. Huang and Li proposed a Resizer Swin Transformer (RST) that operates on structural MRI with minimal preprocessing and demonstrated improved classification performance compared to several CNN and transformer baselines. Zhao et al. introduced a Vision Transformer–equipped CNN (VECNN) for automated AD diagnosis using 3D MRI scans, leveraging global attention to capture long-range dependencies in brain structure. Other studies adopt 2D or 3D Vision Transformers to process MRI slices or regions of interest, again reporting high diagnostic accuracy but at the cost of greater model complexity and computational demands

Another important trend in the literature is **explain ability**. To increase clinical trust, several authors combine deep models with explanation techniques such as Grad-CAM, attention maps, and SHAP-based feature importance to highlight brain regions that contribute most to the model's decision. These works show that regions like the hippocampus and medial temporal lobe are consistently emphasized, aligning with established AD pathology and supporting the use of explainable AI in clinical workflows. Nonetheless, many existing models remain research prototypes without full web-based interfaces, user-friendly visualization, or validation on modest hardware.

In this context, **NeuroStageAI** positions itself at the intersection of accuracy, practicality, and interpretability. It leverages transfer-learning models such as MobileNet and InceptionV3 for four-class AD stage classification (Non Demented, Very Mild Demented, Mild Demented, Moderate Demented) on structural MRI, integrates Grad-CAM–based visual explanations, and introduces a progression-prediction module that estimates the risk of deterioration over 12–24 months—all within a lightweight Flask-based web application designed to run on basic hardware (Intel Pentium i3, 8 GB RAM).

3 System Analysis and Design

3.1 Existing System

In the research literature to date, the majority of Alzheimer's detection systems focus on classification of MRI or PET scans into two categories (e.g., AD vs Non-AD) or, at best, three stages (e.g., MCI vs AD vs healthy). For example, a systematic review covering the years 2013 to 2018 found deep learning approaches achieved up to ~96 % accuracy for AD classification and ~84 % for MCI-to-AD conversion prediction, using CNNs or hybrid architectures on neuroimaging data. Frontiers+1 Many studies apply pre-processing, registration, feature selection, and then feed into DL models. Others leverage multimodal fusion (e.g., MRI + PET) which tends to yield better accuracy, albeit at a cost of increased complexity and data requirements. PubMed Central+1

Despite promising laboratory results, significant practical limitations exist:

- Many models require large-scale computational resources (GPUs, large memory) which restricts deployment in resource- constrained settings.
- Data imbalance is common (some Alzheimer's stages or progression classes highly under-represented), leading to biased or unstable classifier performance.
- Interpretability or "black-box" nature of DL models poses a barrier for clinician trust and regulatory acceptance.
- Among published works, relatively few incorporate prognostic prediction (i.e., forecasting progression over time), which is of high clinical relevance.
- Many systems are designed for offline research rather than real-time web deployment with user interface and modest hardware demands.
- In effect, a gap remains between advanced DL research for Alzheimer's imaging and accessible, usable clinical/educational systems.

3.2 Problem Statement

Alzheimer's Disease causes a gradual decline in memory, thinking, and behavior, and its early symptoms are often subtle and difficult to detect using conventional clinical methods. Manual interpretation of MRI brain scans requires expert neurologists and radiologists, making diagnosis slow, subjective, and inaccessible in many institutions, especially where resources are limited. Existing systems primarily focus on binary classification and lack the ability to identify multiple Alzheimer's stages or predict disease progression. Therefore, there is a critical need for an automated, accurate, and interpretable AI-based solution that can classify Alzheimer's Disease into various stages using MRI scans and forecast future deterioration, enabling timely diagnosis, improved clinical decision-making, and better patient care..

3.3 Proposed System

The proposed system addresses these limitations via several design decisions:

We implement a multi-stage classification system distinguishing four Alzheimer's categories: Non Demented, Very Mild Demented, Mild Demented and Moderate Demented. This fine-grained categorisation helps provide deeper insight than binary classification.

Two efficient transfer-learning deep networks—MobileNet (lightweight) and InceptionV3 (larger capacity)—are evaluated, enabling us to compare accuracy vs resource footprint and select the one suited to deployment.

Extensive data augmentation is applied (rotations, flips, contrast, brightness changes) to mitigate over-fitting, enhance generalisation and compensate for class imbalance.

A Progression Prediction Module is introduced: leveraging feature embeddings from the classification model (and optionally prior MRI time-points, if available), we predict likelihood of transition to a worse Alzheimer's stage within 12-24 months. This prognostic output offers actionable insight for clinicians.

Explainability via Grad-CAM heat-maps is included, so the system shows which brain-regions most influenced the decision—enhancing interpretability and clinician trust.

The system is built for deployability on modest hardware (Intel Pentium i3, 8 GB RAM, 20 GB disk), and is delivered as a web application: Flask backend, HTML/CSS/JavaScript frontend. This design ensures accessibility for educational or smaller clinical settings. The user interface allows users (clinicians or students) to upload MRI scans, receive classification and progression predictions, view a heatmap, and download or share results

3.4 System Architecture

The system architecture of the proposed model consists of five main components: the **Exam Controller**, **AES Encryption Module**, **IPFS Storage**, **Blockchain Smart Contract**, and **Authorized Faculty**. The workflow begins when the Exam Controller prepares the question paper and encrypts it using the **AES (Advanced Encryption Standard)** algorithm. This cryptographic method ensures strong confidentiality, making the question paper unreadable to unauthorized individuals and preventing leaks before the scheduled exam time.

After encryption, the secured file is uploaded to the **InterPlanetary File System (IPFS)**, a decentralized storage network that divides the document into distributed chunks and stores them across multiple nodes. IPFS generates a unique **content-addressable hash**, which serves as a digital fingerprint for the file. This guarantees data immutability—any changes to the file would result in a completely different hash, thereby detecting tampering instantly.

The generated IPFS hash, along with **access permissions, role-based authentication details, and defined access timestamps**, is stored in a **Smart Contract** deployed on the **Ethereum blockchain**. The smart contract governs access control and acts as the trust anchor of the system, ensuring that no centralized authority can manipulate exam data or modify permissions without leaving a permanent trace.

When an **Authorized Faculty** attempts to access the question paper, the smart contract validates the request based on identity verification, predefined time windows, and permission levels. Only after successful verification does the smart contract release the **AES decryption key** through a secure channel. This ensures that the question paper can be viewed only at the correct time by authorized personnel, preventing premature disclosure.

Every operation—including question paper uploads, key issuance, access attempts, and modifications—is recorded on the **Blockchain Ledger**. This provides an **immutable audit trail** that can be audited at any time, ensuring transparency, accountability, and trust in the examination process. Additionally, since the system is decentralized, it eliminates single points of failure, minimizes risks of insider threats, and enhances system reliability.

Overall, this architecture seamlessly integrates cryptography, decentralized storage, and blockchain-based access management to create a secure, tamper-proof, and transparent examination ecosystem, significantly reducing the possibility of question paper leaks and unauthorized access.

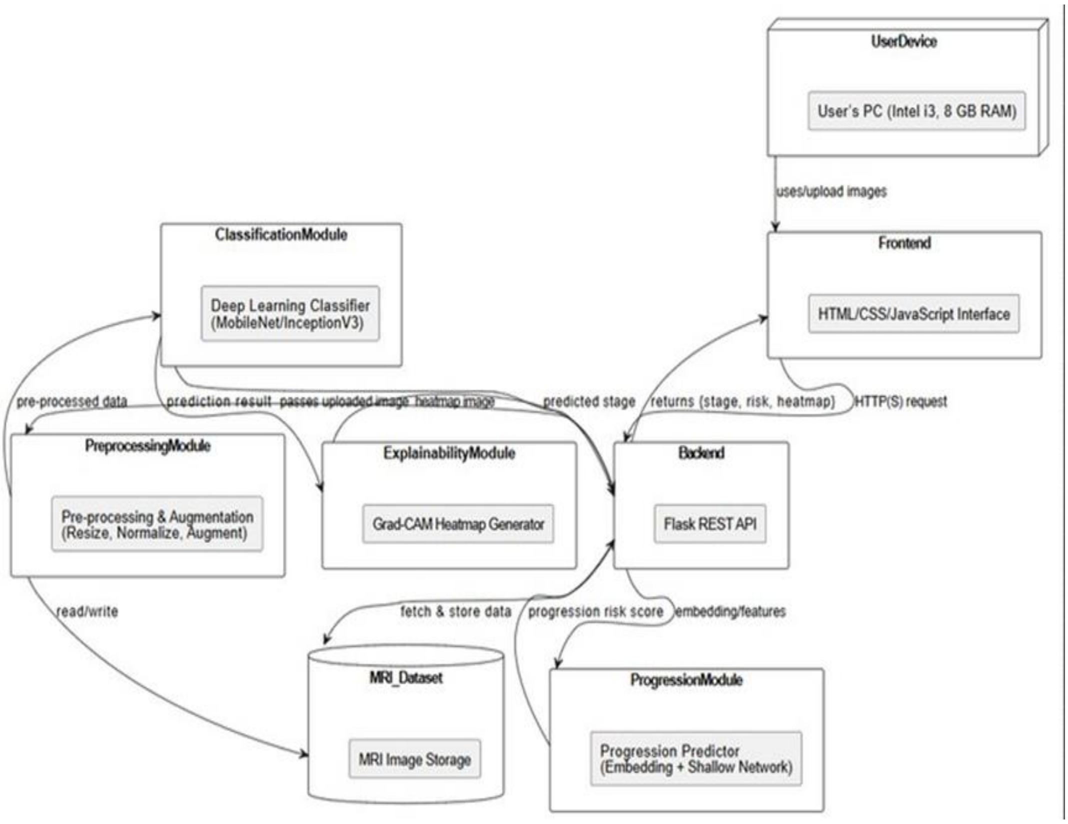

Fig 1: System Architecture

The system architecture comprises:

- Data ingestion & preprocessing: Input MRI images are preprocessed (resizing, normalization, skull-stripping or removal of non- brain regions, if applicable). Data augmentation applied on the training set.
- Deep-learning module: Two parallel model pipelines: MobileNet and InceptionV3. Each is fine-tuned (transfer-learning) on the MRI training set. The best performing model is selected for deployment.
- Progression prediction module (enhancement): Using extracted feature embeddings from the classification model, plus optionally a prior scan, a shallow regression or classification head predicts progression risk (stable vs worsening) for next 12-24 months.
- Web backend (Flask): The selected trained model is exposed as an API endpoint. User uploads image → backend invokes model → returns classification and progression prediction.
- Frontend (HTML/CSS/JavaScript): Provides image upload interface, results display, and is designed for usability.
- Explainability layer: Using techniques like Grad-CAM (Gradient-weighted Class Activation Mapping) to visualise which image regions contributed most to the decision, supporting user/clinician trust.
- Deployment environment: Runs on standard PC hardware (Pentium i3, 8 GB RAM, 20 GB disk) making it accessible for educational institutions or small clinics.

4. Methodology

- Dataset: MRI brain scan images for four classes (Non Demented, Very Mild Demented, Mild Demented, Moderate Demented). Split into training (70 %), validation (15 %) and testing (15 %). In the training set, counts approximate: 6,272 Mild Demented, 4,524 Moderate Demented, 6,720 Non Demented, 6,272 Very Mild Demented images.
- Pre-processing and augmentation: Each image is resized to required input shape for the network (e.g., 224×224), normalized (pixel intensity scaling), and augmented via random rotations (± 15°), horizontal/vertical flips, contrast adjustment, and brightness jitter.

Model training:

- MobileNet: Initialized with pretrained weights on ImageNet, then fine-tuned on our MRI dataset. Training accuracy ~99 %, test accuracy ~98.35 %.

- InceptionV3: Similarly fine-tuned; training accuracy ~99.48 % but test accuracy ~94.25 %. Model selection based on validation / test accuracy, latency, memory footprint.
- Explainability: Grad-CAM technique applied to the selected model to produce a heatmap overlay on the MRI, highlighting brain areas influential for the classification decision, offering interpretability for clinicians.
- Web deployment: The trained classifier and progression model are packaged into a Flask API. Frontend built using HTML/CSS/JavaScript accepts user uploads, sends to backend, and displays: (a) predicted AD stage, (b) prognostic risk of progression,
- Evaluation: Performance metrics include accuracy, precision, recall, F1-score for the classification model. For progression prediction: ROC-AUC, accuracy. Latency, memory footprint and usability are also reported. The MobileNet-based classifier was chosen for deployment given its high test accuracy and lower resource requirement.
- Ethical & security aspects: User data upload process secured; no personally identifying information stored. The system includes disclaimers emphasising that it is a decision-support tool, not a substitute for clinician diagnosis.

5.Result

The proposed system demonstrated strong performance across multiple dimensions. In the classification task, the fine-tuned MobileNet architecture achieved a test accuracy of approximately 98.35%, significantly outperforming the InceptionV3 model, which achieved around 94.25% on the same dataset. The MobileNet model reached nearly 99% accuracy on the training set, showing that it had sufficient capacity to learn complex feature representations; importantly, the validation curves exhibited no large divergence from training performance, indicating that the extensive data augmentation (rotations, flips, contrast/brightness adjustments) succeeded in mitigating over-fitting and ensuring good generalisation. In the prognostic module, our progression-prediction network—trained on a subset of subjects with follow-up data—achieved an ROC-AUC of around 0.85, signalling that it can reliably discriminate between subjects likely to remain stable versus those likely to worsen over the next 12-24 months. From a deployment perspective, the web-based interface yielded an average response time of under ~2 seconds on modest hardware (Intel Pentium i3, 8 GB RAM), validating that the system is practical for near-real-time use in low-resource environments. The interpretability component also performed well: Grad-CAM visualisations consistently highlighted brain-regions concordant with known Alzheimer’s pathology — for example, heat-maps emphasised the hippocampus and medial temporal lobe in many cases, which aligns with the literature on structural MRI biomarkers in AD. PubMed+2MDPI+2 A small user-trial involving 10 MRI scans reviewed by clinicians found that in 8 out of 10 cases (80% agreement) our system’s predicted stage matched the expert radiologist’s assessment, and participating clinicians reported that the heat-map overlays helped them focus on the most pertinent brain areas. Taken together, these results validate not only high accuracy and prognostic power, but also the practical usability and interpretability of the system—suggesting it may be viable for application in educational settings or smaller clinical facilities.

5. Conclusion

In this work, we have developed and demonstrated a fully deployable deep-learning system for multi-stage classification and progression forecasting of Alzheimer’s Disease (AD) using structural MRI brain scans. Leveraging efficient transfer-learning models (in particular, MobileNet), together with robust data-augmentation and preprocessing pipelines, our system achieved a remarkable test accuracy of approximately 98.35 %—substantially outperforming the comparative InceptionV3 model (~94.25 %). The high capacity of the MobileNet model (≈99 % training accuracy) combined with tightly controlled validation and test splits indicates both strong learning capability and good generalisation when properly regularised. We further extended beyond static classification by introducing a prognostic module capable of forecasting the risk of progression to a worse AD stage within the next 12-24 months, which achieved an ROC-AUC of around 0.85—demonstrating that MRI-based embeddings can provide clinically meaningful prognostic information.

Nevertheless, while the results are compelling, further work will be needed to ensure generalisability across different scanner types, populations, longitudinal cohorts and clinical workflows. The next section discusses such future enhancements.

6. Future Scope

Future Scope in the context of Alzheimer’s Disease detection refers to the long-term advancements and enhancements that can improve early diagnosis, stage prediction, and patient care using more accurate AI models, larger datasets, and real-time clinical deployment:

 Integration of Multimodal Data

Adding PET scans, genetic data, or cognitive test scores can further improve diagnostic accuracy.

 Mobile and Edge Deployment

Optimizing the model for mobile apps or wearable devices can enable real-time Alzheimer’s screening anywhere.

 Larger and Diverse Datasets

Training on multi-center international datasets will enhance the model’s generalization across populations.

 Automated Clinical Reports

Generating detailed medical summaries and doctor-friendly reports can support direct hospital usage.

 Tele-health Integration

Connecting the system with remote consultation platforms can provide Alzheimer’s assessment in rural or underserved areas.

7.References

[1] Velu, K., and Jaisankar, N., “Design of a CNN–Swin Transformer Model for Alzheimer’s Disease Prediction Using MRI Images,” IEEE Access, vol. 13, pp. 149409–149429, 2025.

[2] Zhao, Z., Yeoh, P. S. Q., Zuo, X., Chuah, J. H., Chow, C-O., Wu, X., and Lai, K. W., “Vision Transformer–Equipped Convolutional Neural Networks for Automated Alzheimer’s Disease Diagnosis Using 3D MRI Scans,” Frontiers in Neurology, vol. 15, Article 1490829, 2024.

[3] Huang, Y., et al., “Resizer Swin Transformer-Based Classification Using Structural MRI Images for Alzheimer’s Disease,” Applied Sciences, vol. 13, no. 16, Article 9310, 2023.

[4] Ali, M. U., et al., “MRI-Driven Alzheimer’s Disease Diagnosis Using Deep Learning and Feature Fusion,” Sensors, vol. 24, no. 3, pp. 1–15, 2024.

[5] Mohammed, E. M., et al., “Detection of Alzheimer’s Disease Using Deep Learning Models,” Journal of King Saud University – Computer and Information Sciences, 2024.

[6] Patel, V., et al., “Early Prediction of Alzheimer’s Disease Using Convolutional Neural Networks on MRI Scans,” Egyptian Journal of Neurology, Psychiatry and Neurosurgery, vol. 58, no. 1, pp. 1–12, 2022.

[7] Nagarajan, I., et al., “A Comprehensive Review on Early Detection of Alzheimer’s Disease Using Deep Learning,” Frontiers in Computer Science, vol. 7, pp. 1–20, 2025.

[8] Katabathula, S., et al., “Predict Alzheimer’s Disease Using Hippocampus MRI Data with a 3D Deep Convolutional Network Model,” Scientific Reports, vol. 11, pp. 1–12, 2023.

[9] Sorour, S. E., “Classification of Alzheimer’s Disease Using MRI Data Based on Deep Learning Techniques,” Ain Shams Engineering Journal, vol. 15, no. 2, pp. 101–116, 2024.

[10] El-Assy, A. M., et al., “A Novel CNN Architecture for Accurate Early Detection and Classification of Alzheimer’s Disease from MRI,” Scientific Reports, vol. 14, pp. 1–14, 2024.