

Early Detection of Alzheimer’s Disease Using Machine Learning

OM PATKAR¹, VEDANT KHORJEKAR², RHITIKESH GAIKWAD³, SALABHA JACOB⁴

^{1,2,3}*Electronics and Computer Science, Shah and Anchor Kutchhi Engineering College, Mumbai, India*

⁴*Professor, ³Electronics and Computer Science, Shah and Anchor Kutchhi Engineering College, Mumbai, India*

Abstract- Early detection of Alzheimer’s disease (AD) is critical for timely intervention and slowing disease progression. Manual analysis of MRI scans is labor-intensive and subject to inter-clinician variability. This paper presents an end-to-end deep learning system for automated classification of Alzheimer’s stages using brain MRI images. A Convolutional Neural Network (CNN) is trained to classify images into four categories: Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented, achieving a test accuracy of 99.84% on 1,276 held-out samples from the Kaggle Alzheimer MRI Dataset (6,376 images total). The system is deployed via a Streamlit web interface for real-time prediction, enhanced with Grad-CAM visualization to highlight discriminative brain regions. Batch processing and timestamped prediction history tracking are also supported. Experimental results demonstrate sub-second inference latency, confirming the system’s potential as a clinical decision-support tool.

Index Terms- Alzheimer’s Disease, CNN, Deep Learning, Grad-CAM, MRI, Streamlit

I. INTRODUCTION

Alzheimer’s disease (AD) is a progressive neurodegenerative disorder characterized by memory loss, cognitive decline, and behavioral impairment. According to the World Health Organization, over 55 million people worldwide currently live with dementia, with AD accounting for approximately 60–70% of cases [1]. As the global population ages, this number is projected to nearly triple by 2050, placing enormous burden on healthcare systems worldwide.

MRI imaging is widely used for detecting structural brain changes associated with AD, such as hippocampal atrophy and cortical thinning. However, manual interpretation is time-consuming, requires

specialized expertise, and is subject to inter-rater variability arising from clinician fatigue and workload.

Recent advances in deep learning, particularly Convolutional Neural Networks (CNNs), have enabled automated feature extraction and classification directly from raw image data, eliminating the need for handcrafted features [2]. CNNs have demonstrated state-of-the-art performance across numerous medical imaging tasks. This work proposes a CNN-based system integrated with a Streamlit web interface for real-time Alzheimer’s stage classification. Grad-CAM visualization is incorporated for explainability, and the system supports batch processing and prediction history tracking, addressing key practical limitations of existing research prototypes.

II. LITERATURE REVIEW

Early computational approaches to Alzheimer’s detection relied on handcrafted feature engineering — texture analysis, voxel-based morphometry, and statistical descriptors — combined with classifiers such as Support Vector Machines (SVM) and Random Forests. While these achieved moderate accuracy on small datasets, they struggled to generalize across diverse imaging protocols [3].

The emergence of deep learning transformed medical image analysis. Litjens et al. [4] demonstrated that CNNs consistently outperform traditional pipelines in classification and segmentation tasks. Studies on AD detection using CNNs on MRI data report accuracies in the range of 80–96%.

Explainability is critical for clinical adoption. Selvaraju et al. introduced Grad-CAM [5], which generates visual explanations by using gradient information flowing into the final convolutional layer to highlight discriminative regions. This technique has been widely validated against clinical anatomy in medical imaging.

However, most existing works are research prototypes that lack practical deployment features such as user interfaces, batch processing, and result history. This project addresses these gaps with a complete deployable system.



Figure 1: Enter Caption

III. PROPOSED SYSTEM

A. System Overview

The proposed system is an end-to-end pipeline for automated Alzheimer's disease stage classification from brain MRI images. It integrates a trained CNN inference engine with a Streamlit web interface to deliver real-time predictions accessible to clinical users. Given an input MRI image, the system produces a stage classification, confidence score, and Grad-CAM heatmap.

B. System Workflow

The workflow proceeds as: (1) user uploads MRI image via Streamlit; (2) image is preprocessed through grayscale conversion, resizing to 128×128 pixels, and normalization; (3) the preprocessed tensor is passed through the CNN; (4) Softmax probabilities yield the predicted class and confidence score; (5) Grad-CAM heatmap is generated and overlaid on the image; (6) result is displayed and logged to prediction history.

C. System Architecture

The system follows a modular architecture with four in-dependent components: (i) Preprocessing Module for image standardization; (ii) CNN Inference Engine for classification; (iii) Explainability Module for Grad-CAM generation; and (iv) Application Layer for user interaction, batch processing, and history management via Streamlit.

IV. METHODOLOGY

A. Dataset

The Kaggle Alzheimer MRI Dataset was used, comprising 6,376 labeled brain MRI images across four classes: Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented. The dataset was split 80/20 into 5,100 training and 1,276 testing samples.

B. Preprocessing

Each image undergoes: (i) grayscale conversion to eliminate color redundancy; (ii) resizing to uniform 128×128 resolution; and (iii) conversion to a normalized PyTorch tensor with pixel values in $[0, 1]$. This pipeline ensures input consistency and accelerates model convergence.

C. CNN Architecture

The CNN comprises three convolutional blocks, each with a Conv2D layer (ReLU activation) followed by 2×2 max pooling. Filter counts increase progressively: 32, 64, 128. The flattened output feeds a fully connected layer (256 neurons, ReLU), a Dropout layer ($p = 0.5$) for regularization, and a Softmax output layer over 4 classes. Training used Adam optimizer ($lr = 0.001$), Cross-Entropy loss, batch size 32, over 10 epochs.

D. Grad-CAM Visualization

Grad-CAM is applied to the final convolutional layer. Gradients of the predicted class score with respect to the final feature maps are globally average-pooled to obtain importance weights. A weighted combination of feature maps — passed through ReLU — produces a heatmap upsampled and overlaid on the input MRI, highlighting regions most influential in the prediction.

V. SYSTEM IMPLEMENTATION

A. Technology Stack

The system is built entirely on an open-source technology stack to ensure accessibility and reproducibility. Python 3.10 serves as the core programming language due to its extensive ecosystem for scientific computing and machine learning. Py-Torch 2.0 is used for defining, training, and running inference on the CNN model, leveraging its dynamic computation graph for flexible experimentation. Streamlit is employed as the web application framework, enabling rapid deployment of an inter-active browser-based interface without requiring frontend de-velopment expertise. PIL and OpenCV handle all image load-ing, format conversion, and preprocessing operations, ensuring MRI scans maintain their structural integrity throughout the pipeline. Prediction results are persisted using Python's built-in CSV module, providing a lightweight and portable storage solution that requires no database infrastructure.

B. User Interface

The Streamlit-based user interface is designed to be in-tuitive and accessible to clinical users with no programming background. It provides a drag-and-drop file uploader sup-ported both JPEG and PNG image formats. Upon upload, the system instantly displays the input MRI image alongside the predicted Alzheimer's stage and its associated confidence score as a percentage. The Grad-CAM heatmap is rendered as a color overlay directly on the original MRI image, allowing clinicians to visually verify which brain regions influenced the model's decision. The entire interface is served as a web ap-plication accessible from any standard browser on any device, requiring no software installation or configuration on the client side.

C. Batch Processing

To support higher-throughput clinical screening work-flows, the system provides a dedicated batch processing mode that allows users to upload and process multiple MRI images in a single session. All uploaded images are passed through the preprocessing pipeline and CNN inference engine sequen-tially, with results aggregated and presented in a structured tab-ular format showing the image

filename, predicted class, and confidence score for each scan. The complete batch results table can be exported as a CSV file at any time, facilitating seamless integration with existing clinical documentation and electronic health record systems for audit and reporting pur-poses.

D. Prediction History

The system maintains a persistent prediction history log that automatically records every inference performed during a session. Each entry captures the exact timestamp of the pre-diction, the original image filename, the predicted Alzheimer's disease stage, and the model's confidence score, providing a comprehensive longitudinal record of all processed scans. Users can interact with the history log directly from within the interface — searching for specific records by filename or date, filtering results by predicted class, downloading the full history as a CSV file for external analysis, or clearing the log entirely when required. This feature is particularly valuable for tracking disease progression across multiple visits for the same patient over time.

VI. RESULTS AND DISCUSSION

A. Experimental Setup

All experiments used Python 3.10 and PyTorch 2.0 on the Kaggle Alzheimer MRI Dataset (6,376 images, 80/20 split: 5,100 train / 1,276 test). Training ran for 10 epochs, Adam optimizer ($\text{lr} = 0.001$), Cross-Entropy loss, batch size 32.

B. Training Convergence

Training loss decreased from 0.6329 (epoch 1) to 0.0217 (epoch 10). The rapid drop at epoch 2 (0.0620) reflects fast early feature adaptation. Minor later fluctuations are at-tributable to stochastic batch sampling and do not affect final convergence.

C. Classification Performance

The model achieved a test accuracy of 99.84% on 1,276 test samples. Only 2 misclassifications occurred (Non-Demented predicted as Mild Demented). Table ?? presents per-class metrics.

D. Inference Performance

Average inference latency is under 1 second per image on CPU. Batch processing of 50 images

completes in under 30 seconds, confirming real-time suitability.

E. Grad-CAM Analysis

Grad-CAM heatmaps confirm the model focuses on hippocampal and temporal lobe regions — anatomically consistent with Alzheimer’s-related structural atrophy — validating that predictions are driven by clinically meaningful features.

F. Comparison with Prior Work

Figure 2: System Workflow Diagram

G. Discussion

The near-perfect accuracy of 99.84% is noteworthy. It is important to acknowledge that the random splitting strategy employed does not guarantee patient-level separation between training and test sets. In MRI datasets where multiple slices from the same subject may be present, correlated samples across splits could inflate reported accuracy. This is a limitation of the current evaluation methodology. Future work will implement subject-level cross-validation for more clinically rigorous performance estimates. Notwithstanding this, the results demonstrate strong representational capacity of the proposed CNN architecture.

VII. CONCLUSION

This paper presented a CNN-based system for early detection and staging of Alzheimer’s disease from brain MRI images. The model achieved 99.84% test accuracy on 1,276 held-out samples from the Kaggle Alzheimer MRI Dataset. The Streamlit interface significantly enhances usability, while Grad-CAM provides clinically meaningful model interpretability. Sub-second inference latency and batch processing confirm suitability as a real-time decision-support tool. Future work will focus on subject-level data splitting, larger dataset diversity, advanced architectures such as Vision Transformers, and formal clinical validation.

VIII. ACKNOWLEDGMENT

The authors gratefully acknowledge the guidance and support of Prof. Salabha Jacob and the Department of Electronics and Computer Science, Shah and

Anchor Kutchhi Engineering College, Mumbai, India.

REFERENCES

- [1] World Health Organization, “Dementia,” WHO Fact Sheet, 2023.
- [2] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, pp. 436–444, 2015.
- [3] S. Kloppel et al., “Automatic classification of MR scans in Alzheimer’s disease,” *Brain*, vol. 131, pp. 681–689, 2008.
- [4] G. Litjens et al., “A survey on deep learning in medical image analysis,” *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
- [5] R. R. Selvaraju et al., “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *Proc. ICCV*, 2017.
- [6] S. Basaia et al., “Automated classification of Alzheimer’s disease and mild cognitive impairment using a single MRI and deep neural networks,” *NeuroImage: Clinical*, vol. 21, 2019.
- [7] M. A. Ebrahimihafez, S. Luo, and R. Chiong, “Deep learning to detect Alzheimer’s disease from neuroimaging,” *Informatics in Medicine Unlocked*, vol. 18, 2020.
- [8] M. Liu et al., “Multi-modality cascaded convolutional neural networks for Alzheimer’s disease diagnosis,” *Neuroinformatics*, vol. 16, pp. 295–308, 2018.
- [9] M. Hon and N. M. Khan, “Towards Alzheimer’s disease classification through transfer learning,” in *Proc. IEEE BIBM*, 2017.
- [10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” in *Proc. NeurIPS*, 2012.