

A Transformer–Rule Fusion Framework for Early Multistage Alzheimer’s Detection from MRI Scans

S. Sankar Ganesh^{1*}, S. Lalith Kumar², Geethika Kudipudi², Guttula Niharika²

¹Associate Professor, ²UG Student, ^{1,2}Department of Computer Science and Engineering (AI&ML),

^{1,2}Kommuri Pratap Reddy Institute of Technology, Ghanpur, Ghatkesar, 501301, Telangana, India.

*Correspondence: S. Sankar Ganesh

ABSTRACT

Alzheimer’s disease is the most common form of dementia among older individuals and poses a major global health challenge, with nearly 10 million new cases reported annually. This progressive neurological disorder leads to gradual neurodegenerative changes, beginning with mild memory impairment and advancing to loss of social interaction and environmental awareness. According to Alzheimer Disease International (ADI), approximately 75% of dementia cases remain undetected worldwide, making early diagnosis particularly difficult. Currently, effective solutions to halt disease progression remain limited due to the lack of reliable diagnostic and treatment methods. In this work, high-quality Magnetic Resonance Imaging (MRI) scans are analyzed using a pretrained Swin Transformer model to extract high-dimensional visual embeddings that capture subtle neuro-structural variations across four stages: Normal, Very Mild, Mild, and Moderate Alzheimer’s disease. These extracted features are evaluated using multiple baseline classifiers, including Generalized Linear Model (GLM), Generalized Learning Vector Quantization (GLVQ), and Perceptron, to establish comparative performance. The proposed Optimal Rule Fit (ORF) classifier achieves superior results by effectively modeling complex, non-linear relationships within the MRI feature space. A Tkinter-based Graphical User Interface (GUI) is developed to facilitate dataset upload, preprocessing, model training, evaluation, and visualization through confusion matrices and Receiver Operating Characteristic (ROC) curves. Additionally, the system integrates a remote prediction server enabling real-time, cross-device MRI-based Alzheimer detection. Explainable Artificial Intelligence (XAI) techniques further enhance interpretability by identifying brain-region atrophy, ventricular enlargement, and disease severity. With an accuracy of 99.43%, the proposed framework demonstrates strong reliability for early diagnosis, remote clinical support, and neuroimaging research.

Keywords: Alzheimer’s disease, Dementia, Magnetic Resonance Imaging (MRI), Neurodegeneration, Early diagnosis, Brain atrophy, Ventricular enlargement, Neuroimaging, Feature extraction, Disease staging.

1. INTRODUCTION

With a rapidly aging society, the number of people with Alzheimer’s disease (AD) is increasing globally. Approximately 9.9 million new cases of dementia are reported annually, which implies that a new patient is diagnosed with the disease every 3.2 s. The number of people with AD is expected to exceed 100 million by 2050. However, an effective medical treatment for the disease is yet to be devised. Neuroimaging has been used for understanding the progression of AD, and considerable progress has been achieved in the use of deep learning for medical imaging in research and clinical medicine [1]. The macroscopic findings of AD have revealed diffuse brain atrophy.

As shown in fig. 1 the classification accuracy of AD versus healthy control using deep learning of MRI is 91.4%, and that of mild cognitive impairment (MCI) versus AD is 70.1% [2]. The economic evaluation of dynamic susceptibility MRI compared with nonenhanced computer-assisted tomography (CT) is USD 479,500 per quality-adjusted life year, whereas the corresponding comparison for single-photon emission computed tomography (SPECT) with CT is better (higher effectiveness and lower cost) [3]. Regional cerebral blood flow (RCBF) is related to brain metabolism; therefore, changes in CBF reflect variations in neuronal metabolism. From the perspective of neuropathology, subjects with very mild AD typically exhibit abnormal metabolic and RCBF patterns, even at the preclinical stage.




Fig. 1: Advancements in DL for early diagnosis of Alzheimer's disease

A decreased RCBF already occurs in individuals with MCI before they transition to AD [4]. A disrupted cerebral perfusion may cause impaired vascular clearance ability, which promotes the deposition of beta-amyloid and neurofibrillary tangles. Clinical studies have revealed that RCBF alterations are involved in AD pathogenesis. Even before the accumulation of beta-amyloid, subjects with high-risk exhibit changes in cerebral blood flow. In addition to clinical manifestations, RCBF SPECT, such as voxel-as-features, has frequently been used by physicians as a diagnostic tool [5].

2. LITERATURE SURVEY

Gasmi K et al. [6] integrated using an adaptive weight selection process informed by the Cuckoo Search optimization algorithm. The procedure commences with the pre-processing of neuroimaging data to guarantee quality and uniformity. Features are subsequently retrieved from the neuroimaging data by utilizing the EfficientNetV2B3 and Inception-ResNetV2 models. The Cuckoo Search algorithm allocates weights to various models dynamically, contingent upon their efficacy in particular diagnostic tasks. The framework achieves balanced usage of the distinct characteristics of both models through the iterative optimization of the weight configuration. This method improves classification accuracy, especially for early-stage Alzheimer's disease.

Hasan ME et al. [7] proposed a cutting-edge, computer-assisted method based on an advanced deep learning algorithm to differentiate between people with varying degrees of dementia, including healthy, very mild dementia, mild dementia, and moderate dementia classes. In this paper, four separate models were developed for classifying different dementia stages: CNNs built from scratch, pre-trained VGG16 with additional convolutional layers, graph convolutional networks (GCNs), and CNN-GCN models. The CNNs were implemented, and then the flattened layer output was fed to the GCN classifier, resulting in the proposed CNN-GCN architecture. A total of 6400 whole-brain magnetic resonance imaging scans were obtained from the Alzheimer's Disease Neuroimaging Initiative database to train and evaluate the proposed methods. They applied the 5-fold cross-validation (CV) technique for all the models.

Pruthviraja D et al. [8] developed a classification model of Alzheimer's disease (AD) as MRI scans as input modality. The multi-classification is used for AD variety and is classified into four stages. To leverage the capabilities of the pre-trained GoogLeNet model, transfer learning is employed. The GoogLeNet model, which is pre-trained for image classification tasks, is fine-tuned for the specific purpose of multi-class AD classification. Through this process, a better accuracy of 98% is achieved. As a result, a local cloud web application for Alzheimer's prediction is developed using the proposed architectures of GoogLeNet. This application enables doctors to remotely check for the presence of AD in patients.

Carcagnì P et al. [9] improved the automatic detection of dementia in MRI brain data. For this purpose, they used an established pipeline that includes the registration, slicing, and classification steps. The contribution of this research was to investigate for the first time, to their knowledge, three current and promising deep convolutional models (ResNet, DenseNet, and EfficientNet) and two transformer-based architectures (MAE and DeiT) for mapping input images to clinical diagnosis. To allow a fair comparison, the experiments were performed on two publicly available datasets (ADNI and OASIS) using multiple benchmarks obtained by changing the number of slices per subject extracted from the available 3D voxels. The experiments showed that very deep ResNet and DenseNet models performed better than the shallow ResNet and VGG versions tested in the literature. It was also found that transformer architectures, and DeiT in particular, produced the best classification results and were more robust to the noise added by increasing the number of slices. A significant improvement in accuracy (up to 7%) was achieved compared to the leading state-of-the-art approaches, paving the way for the use of CAD approaches in real-world applications.

Lien W-C et al. [10] focused on overcoming this problem. In this study, the detection performance of five CNN models (MobileNet V2 and NASNetMobile (lightweight models); VGG16, Inception V3, and ResNet (heavier weight models)) on medical images was compared to establish a classification model for epidemiological research. Brain scan image data were collected from 99 subjects, and 4711 images were used. Demographic data were compared using the chi-squared test and one-way analysis of variance with Bonferroni's post hoc test. Accuracy and loss functions were used to evaluate the performance of CNN models. The cognitive abilities screening instrument and mini mental state exam scores of subjects with a clinical dementia rating (CDR) of 2 were considerably lower than those of subjects with a CDR of 1 or 0.5. They analyzed the classification performance of various CNN models for medical images and proved the effectiveness of transfer learning in identifying the mild cognitive impairment, mild AD, and moderate AD scoring based on SPECT images.

Agarwal D et al. [11] focused on pre-processing in four steps: N4 bias field correction, denoising, brain extraction, and registration. End-to-end-learning-based deep CNNs were used to discern between different phases of AD. Eight CNN-based architectures were implemented and assessed. The DenseNet264 excelled in both types of classification, with 82.5% accuracy and 87.63% AUC for training and 81.03% accuracy for testing relating to the sMCI vs. AD and 100% accuracy and 100% AUC for training and 99.56% accuracy for testing relating to the AD vs. CN. Deep learning approaches based on CNN and end-to-end learning offer a strong tool for examining minute but complex properties in MR images which could aid in the early detection and prediction of Alzheimer's disease in clinical settings.

Alqahtani S et al. [12] developed deep learning models that focus on early detection and classifying each case, non-demented, moderate-demented, mild-demented, and very-mild-demented, accordingly through transfer learning (TL); an AlexNet, ResNet-50, GoogleNet (InceptionV3), and SqueezeNet by utilizing magnetic resonance images (MRI) and the use of image augmentation. The acquired images, a total of 12,800 images and four classifications, had to go through a pre-processing phase to be balanced and fit the criteria of each model. Each of these proposed models split the data into 80% training and 20% testing. AlexNet performed an average accuracy of 98.05%, GoogleNet (InceptionV3) performed an average accuracy of 97.80%, and ResNet-50 had an average performing accuracy of 91.11%. The transfer learning approach assists when there is not adequate data to train a network from the start, which aids in tackling one of the major challenges faced when working with deep learning. Mandal PK et al. [13] proposed a deep CNN architecture consisting of 7,866,819 parameters. This model comprises three different convolutional branches, each having a different length. Each branch is comprised of different kernel sizes. This model can predict whether a patient is non-demented, mild-demented, or moderately demented with a 99.05% three-class accuracy. In summary, the deep CNN

model demonstrated exceptional accuracy in the early diagnosis of AD, offering a significant advancement in the field and the potential to improve patient care.

Maqsood M et al. [14] proposed system works on an efficient technique of utilizing transfer learning to classify the images by fine-tuning a pre-trained convolutional network, AlexNet. The architecture is trained and tested over the pre-processed segmented (Grey Matter, White Matter, and Cerebral Spinal Fluid) and un-segmented images for both binary and multi-class classification. The performance of the proposed system is evaluated over Open Access Series of Imaging Studies (OASIS) dataset. The algorithm showed promising results by giving the best overall accuracy of 92.85% for multi-class classification of un-segmented images. Odusami M et al. [15] attempted to elucidate this issue with randomized concatenated deep features obtained from two pre-trained models, which simultaneously learn deep features from brain functional networks from MRI images. They experimented with ResNet18 and DenseNet201 to perform the task of AD multiclass classification. A gradient class activation map was used to mark the discriminating region of the image for the proposed model prediction. Accuracy, precision, and recall were used to assess the performance of the proposed system. The experimental analysis showed that the proposed model was able to achieve 98.86% accuracy, 98.94% precision, and 98.89% recall in multiclass classification. The findings indicate that advanced deep learning with MRI images can be used to classify and predict neurodegenerative brain diseases such as AD.

3. PROPOSED SYSTEM

The proposed system is a robust, automated framework for the early detection of Alzheimer's Disease from structural MRI scans, designed around a secure client-server architecture. As shown in fig. 2 it utilizes a Tkinter-based administrative interface for secure data management, where a Swin Transformer replaces traditional CNN to extract high-dimensional, context-rich feature embeddings that capture both local tissue textures and global brain structures. These features power a proposed ORFC, which is trained and benchmarked against baseline models like GLM and GLVQ to demonstrate superior classification accuracy. For real-world application, the trained system is deployed via a Flask web server that integrates an Explainable AI (XAI) layer powered by the Gemini model; this unique addition not only predicts the disease stage but also validates the input image and generates clinically interpretable reports detailing specific indicators like hippocampal atrophy and ventricle enlargement.

1. Secure System Initialization and Data Management: The proposed system begins with a secure administrative layer designed to manage the sensitive medical dataset. An administrator accesses the system via a Tkinter-based Graphical User Interface (GUI), which is protected by a lightweight, high-performance LMDB (Lightning Memory-Mapped Database) authentication system. Once authenticated, the administrator uploads the structural MRI dataset. The system is designed to be agnostic to the specific folder structure, automatically scanning the directory to identify class labels (e.g., Non-Demented, Mild Demented, Moderate Demented) based on subfolder names. This step ensures that the training data is correctly ingested and organized for the deep learning pipeline.




Fig. 2: Proposed system architecture

2. Deep Feature Extraction using Swin Transformer: The core innovation of this system lies in its move away from traditional CNN toward the state-of-the-art Swin Transformer (Hierarchical Vision Transformer). Once the images are pre-processed (resized to 224x224 pixels and normalized), they are passed through the Swin Transformer model (swin_base_patch4_window7_224). Unlike standard CNNs that process images in a fixed grid, the Swin Transformer uses shifted windows to capture both local details (like tissue texture) and global dependencies (like the overall shape of the brain) simultaneously. The model is used here as a feature extractor, converting every MRI image into a highly detailed 768-dimensional numerical feature vector (embedding), which effectively encapsulates the complex patterns associated with Alzheimer's disease.

3. Model Training and Classification: Once the features are extracted, the system splits the data into training and testing sets (maintaining class balance via stratification) to prepare for model building. This phase involves training multiple classifiers to benchmark performance, culminating in the proposed model. The training process is divided into the following internal steps:

- **3.1. Baseline Model 1: Generalized Linear Model (GLM):** The system first trains a GLM. This model serves as a standard statistical baseline. It attempts to find a linear relationship between the input features (the 768-D vectors) and the target Alzheimer's classes. While computationally fast, GLMs often struggle with the complex, non-linear patterns found in medical image data, providing a basic accuracy benchmark that the proposed model must beat.
- **3.2. Baseline Model 2: Perceptron:** Next, the system trains a Perceptron, the simplest type of artificial neural network. This linear binary classifier is used to test if the Alzheimer's features can be separated by a simple straight boundary (hyperplane). Its inclusion highlights the complexity of the dataset; since MRI features are rarely linearly separable, the Perceptron usually yields lower accuracy, demonstrating the need for more advanced, non-linear classification methods.
- **3.3. Baseline Model 3: Generalized Learning Vector Quantization (GLVQ):** The third baseline is GLVQ, a prototype-based classification algorithm. Unlike the previous two, GLVQ classifies data by learning "prototype" vectors that represent each class (e.g., a prototype for 'Mild Demented'). It classifies new images based on which prototype they are closest to. This is a powerful distance-based method but can be sensitive to the high dimensionality (768 dimensions) of the Transformer output, sometimes leading to overfitting or "curse of dimensionality" issues.

- 3.4. Proposed Model: ORFC:** Finally, the system trains the Proposed ORFC. The ORFC is designed specifically to overcome the limitations of the three baseline models. Instead of relying on simple linear boundaries (like GLM/Perceptron) or pure distance metrics (like GLVQ), the ORFC identifies specific "optimal regions" in the high-dimensional feature space that strongly correlate with each disease stage. It aggregates weak learners (similar to random forests) to create a robust, non-linear decision boundary. This model is designated as the "Proposed" method because it effectively handles the complex, high-dimensional embeddings from the Swin Transformer, yielding the highest accuracy and F1-score among all tested models.

4. Deployment via Flask (Client-Server Architecture): To make the system practical for real-world medical scenarios, the architecture is decoupled into a training environment and a prediction server. The trained components (the Swin Transformer configuration and the saved ORFC model) are deployed on a separate machine running a Flask Web Server. This server exposes a REST API endpoint (/predict) that listens for incoming MRI images from remote clients. This design allows the heavy computational task of training to happen offline on a powerful machine (Laptop 1), while the prediction service runs efficiently on a deployed server (Laptop 2), ready to serve doctors or diagnostic tools in real-time.

5. Explainable AI (XAI) Integration and Final Prediction: The final step ensures prediction reliability and interpretability. When the Flask server receives an image, it does not immediately classify it. First, it sends the image to an XAI module powered by the Gemini Large Multimodal Model. This XAI layer performs two critical tasks: first, it validates that the uploaded image is indeed a Brain MRI (rejecting non-medical images); second, it extracts human-readable symptoms such as "Hippocampal Atrophy" or "Ventricle Enlargement." Once validated, the image is processed by the local Swin Transformer and ORFC model to generate the final classification (e.g., "Very Mild Demented"). The system returns a comprehensive JSON response containing both the technical classification and the medical explanation, bridging the gap between "Black Box" AI and clinical interpretation.

3.2 Swin Transformer Feature Extraction

The Swin Transformer, specifically the Base model initialized with pre-trained weights, serves as the non-traditional, hierarchical feature extraction backbone for this Alzheimer's detection system, directly replacing a standard CNN. Its role is highly specific: by removing its original classification head, the model is utilized solely to process the preprocessed (224x224) MRI images and output a concise 768-dimensional feature embedding. As shown in fig. 3 the process leverages the Swin Transformer's shifted window mechanism to efficiently capture both fine-grained local structural anomalies (like early atrophy) and broad global context across the entire brain slice, providing a superior, highly descriptive input vector for the subsequent, lighter-weight ORFC, thereby improving diagnostic accuracy and efficiency over conventional deep learning models.

1. Initialization and Weights Loading

The process begins by initializing the Swin Transformer Base model (swin_base_patch4_window7_224) using the timm library in PyTorch. Crucially, the model is initialized with pre-trained weights (likely from the large-scale ImageNet dataset) and configured with num_classes=0. This configuration is essential: it removes the final classification head, preventing the model from predicting ImageNet classes and instead forcing it to output the raw feature vector from its final pooling layer. The entire model is transferred to the available device, preferably GPU (cuda), for efficient processing.

2. Image Preprocessing and Scaling

Before the MRI slice can be fed into the Swin Transformer, it must be standardized to meet the model's input requirements. This is handled by a defined transform. Compose pipeline:

- The raw MRI image is resized to a fixed input resolution of (224, 224) pixels.

- It is converted to a PyTorch Tensor and normalized using a specific mean and standard deviation (0.5 for all channels) which matches the standard pre-trained setup. This standardization ensures the input data falls within the optimal range the pre-trained weights were designed for.




Fig. 3: Internal workflow of swin Transformer.

3. Feature Extraction (Forward Pass)

The pre-processed image tensor is passed through the initialized Swin Transformer. The model performs its hierarchical processing, where feature maps are generated across multiple stages, gradually capturing features from local patches to global spatial relationships:

- Patch Partitioning:** The input image is first split into non-overlapping patches.
- Shifted Window Attention:** This mechanism is the core innovation, calculating self-attention within local "windows" and then shifting the windows between layers. This allows the model to capture long-range dependencies (global context, crucial for identifying overall atrophy) while maintaining local details (like tissue density), which is often superior to the fixed local kernels used in traditional CNNs.
- Vector Output:** The final stage aggregates the high-level features into a single 768-dimensional feature embedding.

4. Preparation for Classification

The resulting 768-dimensional feature tensor is moved from the GPU back to the CPU, converted into a NumPy array, and flattened (though the output is already a vector, this step ensures consistent shape (1, 768)). This vector is the comprehensive representation of the original MRI slice. It is this final, information-rich feature vector that is used as the input for the subsequent classification layer, specifically the Proposed ORFC, to make the final diagnosis of the Alzheimer's stage. This separation ensures the Transformer specializes in feature representation, while the ORFC specializes in distinguishing between the classes.

4. RESULTS ANALYSIS

The results section presents the key findings of the study in a clear and organized manner. It highlights the main outcomes derived from the data analysis, showing patterns, trends, or relationships relevant to the research objectives. The information is often supported by tables, charts, or graphs for better understanding. This section focuses only on factual findings without interpretation. It helps readers quickly grasp what was discovered during the study. It forms the foundation for the discussion and conclusions that follow.




Fig. 4: Confusion matrix obtained using ORFC

Fig. 4 shows the confusion matrix demonstrates the classification performance of the Proposed ORFC across the Alzheimer MRI categories Mild, Moderate, Normal, and Very Mild using Swin Transformer embeddings. The diagonal dominance shows near-perfect classification for Mild, Moderate, and Normal, where almost all samples fall into the correct predicted category. Very Mild cases also achieve extremely high accuracy, with only a few misclassified as Normal, indicating strong separation capability even for subtle early-stage dementia variations. Compared to baseline models, ORFC sharply reduces inter-class confusion and handles complex MRI feature patterns with far greater robustness. The overall distribution highlights ORFC's superior ability to map high-dimensional transformer features to distinct Alzheimer stages, making it the most reliable classifier in the system.




Fig. 5: ROC Curve obtained using ORFC

Fig. 5 ROC curve plot shows the one-vs-rest performance of the Proposed ORFC classifier, demonstrating exceptional discriminative ability across all Alzheimer MRI categories Mild, Moderate, Normal, and Very Mild. Each class achieves an AUC of 1.00, indicating perfect separation between true positives and false positives, even for subtle early-stage variations such as Very Mild dementia. The curves from all classes almost overlap at the top-left corner of the graph, reflecting extremely high sensitivity with near-zero false-positive rates. The micro-average AUC of 1.00 further confirms flawless multi-class classification performance when aggregating predictions across all categories. Compared to the diagonal random-guess baseline, ORFC shows a dramatic improvement, validating its superior capability to leverage Swin Transformer features for highly accurate Alzheimer stage identification.




Fig. 6: XAI-Supported MRI Interpretation with Transformer-Based Alzheimer Classification

Fig. 6 shows the output panel's complete prediction workflow, showing the original MRI image, the XAI-generated Alzheimer analysis, and the final classification result from the Swin-Transformer with ORFC model. The left section displays the raw brain MRI provided by the user, while the center text block summarizes the XAI engine's interpretation, identifying key biomarkers such as cortical atrophy, ventricular enlargement, and suspected Alzheimer severity. On the right, the transformer classifier overlays its predicted stage "Very Mild" directly onto the processed MRI, highlighting the system's ability to combine deep-learning-based classification with transparent, interpretable medical insights. This unified view allows clinicians and users to visually verify the prediction alongside meaningful neuro-structural explanations.

4.1 Comparative Analysis

Comparative analysis plays an important role in evaluating the effectiveness of different machine learning models used in medical image classification. In this study, several classification algorithms are implemented to analyze the extracted MRI image features and identify patterns associated with Alzheimer's disease. Each model has different learning capabilities, prediction mechanisms, and performance characteristics. By comparing multiple models using standard evaluation metrics such as accuracy, precision, recall, and F1-score, it becomes possible to determine which algorithm provides better classification results. This analysis helps in understanding the strengths and limitations of each model when applied to MRI image data. The comparative evaluation also supports the selection of an efficient model for reliable Alzheimer's disease detection.

The performance comparison table 1 summarizes the effectiveness of four different classifiers GLM, GLVQ, Perceptron, and the Proposed ORF model when applied to Alzheimer stage prediction using Swin Transformer feature embeddings. The GLM classifier demonstrates strong and balanced performance with an accuracy of 93.93% and similar precision, recall, and F-score values, indicating reliable but not optimal separability. GLVQ performs noticeably lower, achieving only 81.25% accuracy, reflecting its limitations in handling complex MRI feature patterns. The Perceptron classifier performs better, reaching 92.25% accuracy, but still falls short in distinguishing subtle early-stage variations. In contrast, the Proposed ORF classifier significantly outperforms all other models, achieving an exceptional 99.43% across accuracy, precision, recall, and F-score, highlighting its superior capability to map high-dimensional transformer-derived features to distinct Alzheimer categories with perfect classification.

Table 1: Performance comparison for the GLM, GLVQ, Perceptron and Proposed ORF Model

Algorithms Name	Accuracy	Precision	Recall	F-score
GLM Classifier	93.93%	94.45%	93.935%	93.89%
GLVQ Classifier	81.25%	80.40%	81.25%	80.68%
Perceptron Classifier	92.25%	92.54%	92.25%	92.19%
ORF Classifier	99.43%	99.43%	99.43%	99.43%

5. CONCLUSION

The research successfully demonstrates a comprehensive and intelligent framework for automated Alzheimer's disease detection and stage classification using structural MRI scans, leveraging cutting-edge deep learning and machine learning techniques. By integrating the Swin Transformer for high-quality feature extraction with multiple classical classifiers and the proposed ORF model, the system effectively captures subtle structural changes in the brain that distinguish different Alzheimer severity levels. Experimental results confirm that while traditional models like GLM, GLVQ, and Perceptron offer reasonable performance, they are limited in handling complex MRI feature distributions; in contrast, the proposed ORF classifier achieves exceptional accuracy of 99.43%, proving its superiority in learning non-linear and fine-grained patterns. The research also incorporates an interactive Tkinter GUI for dataset management, feature extraction, model training, evaluation, and visualization, making it accessible to non-technical users. Furthermore, a Flask-based remote prediction server enables cross-device inference, allowing MRI images to be uploaded from another laptop or clinic system for real-time Alzheimer stage prediction. The system is enriched with XAI-supported interpretability, providing clinically relevant insights such as brain-region atrophy, ventricle enlargement, and severity estimation, thereby increasing transparency and trust in AI-driven diagnosis.

REFERENCES

- [1] Chetelat, G.A.; Baron, J.-C. Early diagnosis of Alzheimer's disease: Contribution of structural neuroimaging. *Neuroimage* 2003, 18, 525–541.

- [2] Hashizume, Y. Macroscopic findings of brain with ementia. *Neuropathology* 2022, 42, 353–366.
- [3] Li, F.; Tran, L.; Thung, K.H.; Ji, S.; Shen, D.; Li, J. A robust deep model for improved classification of AD/MCI patients. *IEEE J. Biomed. Health Inform.* 2015, 19, 1610–1616.
- [4] Leung, G.M.; Yeung, R.Y.; Chi, I.; Chu, L.W. The economics of Alzheimer disease. *Dement. Geriatr. Cogn. Disord.* 2003, 15, 34–43.
- [5] Lien, W.-C.; Wang, W.-M.; Wang, F.; Wang, J.-D. Savings of loss-of-life expectancy and lifetime medical costs from prevention of spinal cord injuries: Analysis of nationwide data followed for 17 years. *Inj. Prev.* 2021, 27, 567–573.
- [6] Gasmi K, Alyami A, Hamid O, Altaieb MO, Shahin OR, Ben Ammar L, Chouaib H, Shehab A. Optimized Hybrid Deep Learning Framework for Early Detection of Alzheimer’s Disease Using Adaptive Weight Selection. *Diagnostics.* 2024; 14(24):2779. <https://doi.org/10.3390/diagnostics14242779>
- [7] Hasan ME, Wagler A. New Convolutional Neural Network and Graph Convolutional Network-Based Architecture for AI Applications in Alzheimer’s Disease and Dementia-Stage Classification. *AI.* 2024; 5(1):342-363. <https://doi.org/10.3390/ai5010017>
- [8] Pruthviraja D, Nagaraju SC, Mudligiriappa N, Raisinghani MS, Khan SB, Alkhalidi NA, Malibari AA. Detection of Alzheimer’s Disease Based on Cloud-Based Deep Learning Paradigm. *Diagnostics.* 2023; 13(16):2687. <https://doi.org/10.3390/diagnostics13162687>
- [9] Carcagnì P, Leo M, Del Coco M, Distante C, De Salve A. Convolution Neural Networks and Self-Attention Learners for Alzheimer Dementia Diagnosis from Brain MRI. *Sensors.* 2023; 23(3):1694. <https://doi.org/10.3390/s23031694>
- [10] Lien W-C, Yeh C-H, Chang C-Y, Chang C-H, Wang W-M, Chen C-H, Lin Y-C. Convolutional Neural Networks to Classify Alzheimer’s Disease Severity Based on SPECT Images: A Comparative Study. *Journal of Clinical Medicine.* 2023; 12(6):2218. <https://doi.org/10.3390/jcm12062218>
- [11] Agarwal D, Berbis MA, Martín-Noguerol T, Luna A, Garcia SCP, de la Torre-Díez I. End-to-End Deep Learning Architectures Using 3D Neuroimaging Biomarkers for Early Alzheimer’s Diagnosis. *Mathematics.* 2022; 10(15):2575. <https://doi.org/10.3390/math10152575>
- [12] Alqahtani S, Alqahtani A, Zohdy MA, Alsulami AA, Ganesan S. Severity Grading and Early Detection of Alzheimer’s Disease through Transfer Learning. *Information.* 2023; 14(12):646. <https://doi.org/10.3390/info14120646>
- [13] Mandal PK, Mahto RV. Deep Multi-Branch CNN Architecture for Early Alzheimer’s Detection from Brain MRIs. *Sensors.* 2023; 23(19):8192. <https://doi.org/10.3390/s23198192>
- [14] Maqsood M, Nazir F, Khan U, Aadil F, Jamal H, Mehmood I, Song O-y. Transfer Learning Assisted Classification and Detection of Alzheimer’s Disease Stages Using 3D MRI Scans. *Sensors.* 2019; 19(11):2645. <https://doi.org/10.3390/s19112645>
- [15] Odusami M, Maskeliūnas R, Damaševičius R. An Intelligent System for Early Recognition of Alzheimer’s Disease Using Neuroimaging. *Sensors.* 2022; 22(3):740. <https://doi.org/10.3390/s22030740>