Archives available at journals.mriindia.com

International Journal on Advanced Computer Theory and Engineering

ISSN: 2319-2526

Volume 15 Issue 01, 2026

Developing a Data-Driven System for the Early Identification of Alzheimer's Disease through MRI Analysis

[1]Karan Hemant Banerjee, [2]Faiz Kafil Moulavi, [3]Dyuti Agarwal, [4]Sneha Prem Choudhary, [5]Natali Sankhe, [6]Manasi Churi
[1-6] Department of Computer Engineering, Shree L.R. Tiwari College of Engineering, Mumbai
Email: [1]karan.h.banerjee@slrtce.in, [2]faiz.k.moulavi@slrtce.in, [3]dyuti.agarwal@slrtce.in, [4]sneha.p.choudhary@slrtce.in, [5]natali.sankhe@slrtce.in, [6]manasi.churi@slrtce.in

Peer Review Information

Submission: 24 March 2026

Revision: 13 April 2026

Acceptance: 27 April 2026

Keywords

Alzheimer's Disease, Deep Learning, MRI Analysis, ResNet-18, Medical Image Classification.

ABSTRACT

Alzheimer's Disease is one of the most common neurological disorders and currently affects more than 55 million people worldwide. Detecting the disease in its early stages is important because early treatment can help slow down the progression of symptoms. In most hospitals, MRI brain scans are examined manually by radiologists, which can take time and may sometimes lead to inconsistent interpretations.

This project presents a deep learning-based system designed to classify MRI brain images into different stages of Alzheimer's disease using the OASIS dataset. One of the main difficulties encountered during this work was the severe imbalance between classes. The Moderate Dementia category represents only a very small portion of the dataset compared to the healthy cases. To address this issue, techniques such as Weighted Random Sampling and Weighted Cross-Entropy Loss were applied during model training.

Several neural network architectures were evaluated, including a basic Convolutional Neural Network (CNN), EfficientNet-B0, and a fine-tuned ResNet-18 model. Among these, ResNet-18 produced the best performance, achieving a test accuracy of 96.82% and a Macro F1-score of 0.96. The model also achieved 100% recall for the Moderate Dementia class, meaning that all severe cases were successfully detected during evaluation.

To demonstrate practical usability, the model was integrated into a web-based system using FastAPI for the backend and Next.js for the frontend, allowing MRI images to be analyzed quickly on consumer-level hardware.

Introduction

Alzheimer's Disease is a progressive neurological disorder and is recognized as the most common cause of dementia worldwide, accounting for nearly 60–80% of dementia cases. According to the World Health Organization [1], dementia affects millions globally and continues to rise as populations age. The disease develops due to abnormal biological processes that occur within the brain. Two key pathological characteristics define Alzheimer's disease: the buildup of amyloid-beta plaques outside neurons and the formation of tau protein tangles within neurons. These abnormal protein deposits disrupt the normal communication between brain cells and gradually lead to the degeneration and death of neurons.

Over time, these microscopic biological changes result in noticeable structural alterations in the brain. Medical imaging studies commonly reveal a reduction in the size of the cerebral cortex, significant shrinkage of the hippocampus, and enlargement of the brain ventricles. These structural variations are widely used as biomarkers in neuroimaging-based diagnosis and have been extensively studied in MRI-based research datasets such as OASIS and ADNI [9], [24].

According to recent global estimates, nearly 60 million people are currently living with dementia, and the associated economic burden exceeds 1.5 trillion dollars annually [1]. This growing concern has led researchers to emphasize early detection, as timely diagnosis significantly improves the effectiveness of therapeutic interventions.

Detecting the condition during the Very Mild or Mild stages increases the chances of slowing disease progression. Modern therapeutic approaches, including monoclonal antibody treatments, are more effective when applied early. However, early diagnosis remains challenging due to subtle structural changes that are difficult to detect manually.

In many clinical environments, diagnosis depends on a combination of patient history, cognitive tests such as the Mini-Mental State Examination (MMSE), and visual inspection of MRI scans. While MRI is a powerful imaging modality, early-stage Alzheimer's changes are often too subtle for consistent manual interpretation. Studies such as Cuingnet et al. [22] have shown that even advanced computational methods can struggle with classification due to overlapping features between normal aging and early disease stages.

Additionally, radiologists may face challenges such as fatigue, subjective interpretation, and variability in diagnosis. This is further compounded by the global shortage of trained specialists, limiting large-scale screening efforts.

To address these limitations, Computer-Aided Diagnosis (CAD) systems based on artificial intelligence have gained significant attention. Deep learning, particularly Convolutional Neural Networks (CNNs), has demonstrated strong performance in medical image analysis [10], [15]. These models can automatically learn hierarchical feature representations from MRI scans, capturing patterns that are often imperceptible to the human eye.

Early successes of deep learning in medical imaging, including applications in skin cancer classification [19] and brain tumor segmentation [20], have demonstrated the potential of such approaches in clinical diagnostics. In Alzheimer's research, similar techniques have been applied to classify disease stages using MRI data [21], [25].

The primary objective of this research is to design a data-driven diagnostic framework capable of classifying MRI brain images into four categories: Non-Demented, Very Mild Dementia, Mild Dementia, and Moderate Dementia. The system leverages deep learning architectures trained on datasets such as OASIS [9].

In addition to accuracy, this study focuses on computational efficiency. Many existing systems rely on high-performance computing infrastructure, making them impractical for real-world deployment. This research explores the feasibility of running the model on consumer-level hardware, making advanced diagnostic tools accessible to smaller healthcare facilities.

Problem Statement

As of 2026, healthcare systems around the world are experiencing increasing difficulty in managing Alzheimer's Disease, which currently affects nearly 60 million people globally. Although recent advancements in the pharmaceutical field have introduced newer monoclonal antibody therapies and treatments targeting neuro-inflammation, the effectiveness of these treatments depends greatly on detecting the disease at an early stage. In particular, medical intervention is most beneficial when Alzheimer's is identified during the Very Mild or Mild stages of cognitive decline. However, despite these medical developments, the diagnostic process still faces several major challenges that limit early detection.

1. Clinical Diagnostic Limitations and Human Subjectivity

In many hospitals, the diagnosis of Alzheimer's disease still depends largely on the manual evaluation of Magnetic Resonance Imaging (MRI) scans performed by radiologists. Although MRI scans provide detailed images of brain structures, the early indicators of Alzheimer's disease such as small reductions in hippocampal volume or mild thinning of the cerebral cortex can be extremely subtle. These changes may appear similar to normal age-related brain atrophy, making it difficult to clearly distinguish early disease patterns through visual inspection alone.

As a result, different radiologists may interpret the same scan differently, leading to variations in diagnosis. In addition, there is a shortage of trained neuroradiologists in many regions of the world. This shortage contributes to what is often referred to as a diagnostic or treatment gap, where a large percentage of dementia cases remain undetected until the disease has already progressed to advanced and irreversible stages. Current estimates suggest that nearly three-quarters of dementia cases worldwide a not diagnosed early enough, which prevents patients from benefiting from newer treatment options.

2. Impact of Dataset Imbalance on Artificial Intelligence Models

Although artificial intelligence has shown promising results in medical image analysis, the development of reliable AI systems for Alzheimer's detection is complicated by the nature of available clinical datasets. Many datasets used in research, including the widely used OASIS MRI dataset, contain a highly uneven distribution of samples across different disease stages.

For example, the Moderate Dementia category represents only a very small percentage of the total data, while Non-Demented images make up the majority of the dataset. This imbalance creates difficulties for standard machine learning algorithms. Since most models are designed to optimize overall accuracy, they tend to favor the majority class during training. As a result, the model may appear to perform well statistically while failing to correctly identify the minority classes.

In medical applications, this type of error is particularly concerning. A model that correctly identifies healthy patients but fails to detect severe dementia cases could lead to serious clinical consequences. Therefore, addressing class imbalance is a critical step in building reliable AI-based diagnostic systems.

3. Computational and Infrastructure Challenges

Another important challenge in applying artificial intelligence to medical imaging is the computational requirement of modern deep learning models. Many state-of-the-art neural networks used for medical image analysis rely on powerful computing infrastructure, including large cloud-based GPU clusters or high-performance computing systems. While these systems can achieve high accuracy, they are often expensive and difficult to maintain.

This creates a technological gap between well-funded research institutions and smaller healthcare facilities. Hospitals in rural areas or regions with limited resources may not have access to the infrastructure needed to deploy such advanced AI systems. As a result, many existing solutions remain limited to research environments rather than being used in real clinical settings.

Because of these limitations, there is a need to develop diagnostic frameworks that can achieve high accuracy while also being computationally efficient. A system that can operate on consumer-level hardware would make AI-based diagnostic tools more accessible to a wider range of medical institutions.

Literature Review

1. Foundations of Deep Learning and CNNs in Neuroimaging

The foundation of modern deep learning was significantly advanced by LeCun, Bengio, and Hinton [2], along with earlier breakthroughs such as AlexNet by Krizhevsky et al. [11], which demonstrated the power of deep convolutional architectures in large-scale image classification. Further theoretical advancements by Goodfellow et al. [14] and Hinton et al. [12] contributed to improving neural network training and generalization.

CNN architectures have since become the backbone of medical image analysis. Surveys by Litjens et al. [10], Shen et al. [15], and Suzuki [18] highlight how deep learning has transformed the field by enabling automated feature extraction and high-accuracy predictions.

In Alzheimer's disease research, Suk et al. [3] proposed early deep learning-based feature representations for classification tasks, while Sarraf and Tofighi [4] introduced CNN-based approaches for MRI and fMRI analysis. Similarly, Islam and Zhang [5] demonstrated effective Alzheimer's diagnosis using deep learning models applied to brain MRI data.

Despite these advances, Wen et al. [23] identified key methodological issues such as data leakage and lack of reproducibility in CNN-based studies. Their work emphasized the importance of subject-level data splitting, which is adopted in the present study.

2. Advanced CNN Architectures and Transfer Learning

The introduction of deeper architectures such as VGGNet [7], ResNet [6], and Inception models [13] significantly improved image classification performance. ResNet, in particular, addressed the vanishing gradient problem through residual connections, enabling deeper and more effective networks.

Transfer learning has emerged as a practical solution in medical imaging due to limited labeled data. Pretrained models on large datasets like ImageNet can be fine-tuned for domain-specific tasks. Studies such as Basaia et al. [21] demonstrated the effectiveness of transfer learning in Alzheimer's classification.

EfficientNet, proposed by Tan and Le [8], introduced a compound scaling method that optimizes model performance while reducing computational cost. This makes it especially suitable for deployment in resource-constrained environments.

3. Specialized Architectures for Medical Imaging

Beyond standard CNNs, specialized architectures have been developed for medical imaging tasks. Fully Convolutional Networks (FCNs) [16] and U-Net [17] have been widely used for segmentation tasks, enabling precise localization of affected brain regions.

3D CNNs, as explored by Payan and Montana [25], extend traditional CNNs to volumetric MRI data, capturing spatial relationships across slices. These approaches have shown improved performance in detecting Alzheimer's-related structural changes.

Additionally, comparative studies such as Cuingnet et al. [22] evaluated multiple machine learning methods on MRI datasets, highlighting the complexity of Alzheimer's classification and the need for robust models.

4. Dataset Challenges and Class Imbalance

Medical imaging datasets such as OASIS [9] often suffer from class imbalance, where healthy samples significantly outnumber diseased cases. This imbalance can bias models toward majority classes, reducing sensitivity for detecting severe conditions.

Wen et al. [23] emphasized the importance of proper evaluation strategies, while recent studies have proposed techniques such as weighted loss functions and data augmentation to improve minority class detection.

Furthermore, the Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset [24] has played a crucial role in advancing research by providing standardized imaging data for benchmarking algorithms.

5. Deep Learning Applications in Clinical Diagnosis

Deep learning has demonstrated near-human or even superior performance in several medical domains. Esteva et al. [19] achieved dermatologist-level accuracy in skin cancer classification, while Havaei et al. [20] developed robust models for brain tumor segmentation.

In Alzheimer's diagnosis, Basaia et al. [21] showed that deep neural networks can effectively distinguish between Alzheimer's disease, mild cognitive impairment, and healthy controls using MRI data.

These advancements highlight the growing potential of AI-based systems in assisting clinicians, improving diagnostic accuracy, and enabling early disease detection.

Methodology

1. Data Collection and Dataset Organization

The OASIS MRI Dataset was selected as the primary data source. This open-access dataset is widely used in neurological research, particularly for studying brain aging and Alzheimer's disease progression. The dataset contains approximately 20,960 axial two-dimensional T1-weighted MRI slices of the brain.

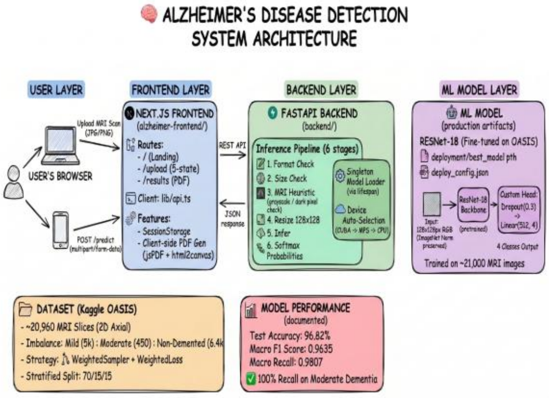

Fig 1: System Architecture of the Proposed Alzheimer's Detection Framework

To maintain reliable model evaluation, the dataset was divided using a stratified sampling strategy. Specifically, 70% of the data was used for model training, while 15% was allocated for validation and the remaining 15% for testing.

Stratification plays an important role in this process because it ensures that all dementia categories, including the minority class "Moderate Dementia," are proportionally represented in each subset. Without stratification, there is a risk that severe dementia cases might not appear in the validation or testing data, which would negatively affect the reliability of model evaluation.

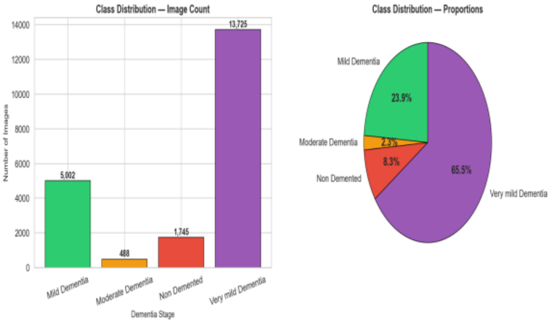

Fig 2: OASIS MRI Dataset Class Distribution

2. Image Preprocessing and Standardization

Raw MRI scans typically contain noise and variations in intensity due to differences in scanning devices and imaging environments. To reduce these inconsistencies and prepare the images for deep learning analysis, a structured preprocessing pipeline was implemented.

First, spatial normalization was performed by resizing all MRI slices to a consistent resolution of 128 × 128 pixels using bilinear interpolation. This resolution was chosen as a compromise between preserving anatomical information and maintaining manageable computational requirements

Next, intensity normalization was applied using Z-score normalization based on commonly used ImageNet statistical parameters ($\mu = [0.485, 0.456, 0.406]$ and $\sigma = [0.229, 0.224, 0.225]$). This step helps align the MRI pixel intensity distribution with the input expectations of pre-trained deep learning models.

Finally, a Non-Local Means (NLM) filtering technique was used to reduce the presence of Rician noise, which is commonly found in T1-weighted MRI scans. This filtering approach removes unwanted artifacts while preserving important structural features such as the hippocampus and cortical regions that are critical for Alzheimer's disease analysis.

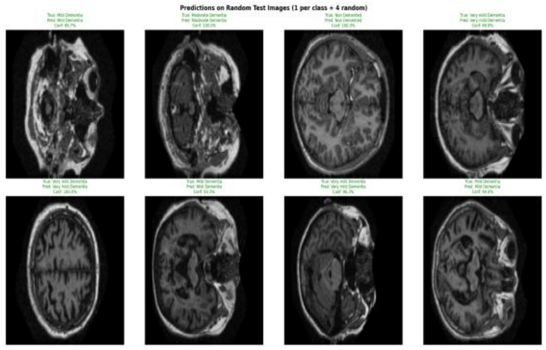

Fig 3: Sample MRI Brain Images

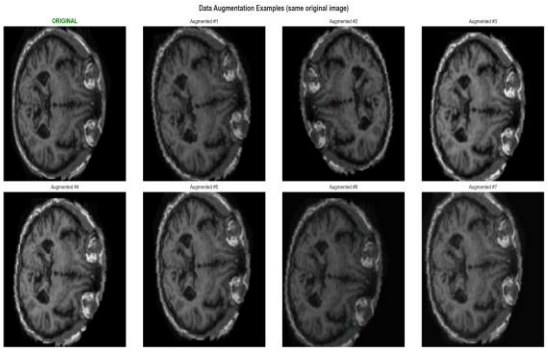

Fig 4: MRI Data Augmentation Examples

3. Handling Dataset Imbalance

One of the major challenges in this dataset is the extreme imbalance between classes. In particular, the "Moderate Dementia" category represents only around 0.5% of the total dataset. If not properly addressed, this imbalance can cause the model to become biased toward the majority classes.

To mitigate this issue, two complementary strategies were adopted. The first approach involves Weighted Random Sampling. Instead of selecting samples sequentially during training, each MRI slice is assigned a sampling probability that is inversely proportional to the frequency of its class.

As a result, samples from underrepresented categories, especially Moderate Dementia, appear more frequently during training batches.The second strategy focuses on cost-sensitive learning. The standard Cross-Entropy Loss function was modified by incorporating class-specific weights. This allows the model to assign a higher penalty when misclassifying samples belonging to rare categories. The loss function for class i is defined as:

$L(i) = -\log((e^\wedge z_i \times w_i) / \Sigma(e^\wedge z_j \times w_j))$

where w_i represents the weight coefficient associated with class i. For the Moderate Dementia category, the weight value was set to 10.73. This adjustment forces the model to pay greater attention to severe dementia cases during optimization.

4. Model Architecture and Training Strategy

The deep learning architecture selected for this study is ResNet-18. This model is well known for its residual learning framework, which uses skip connections to allow gradients to propagate efficiently through deeper network layers. These residual connections help prevent the vanishing gradient problem, which can otherwise hinder the training of deep neural networks, particularly in medical imaging tasks.

To adapt the architecture for Alzheimer's disease classification, the original ImageNet output layer was replaced with a custom classification head. This head consists of a dropout layer with a probability of 30%, followed by a fully connected dense layer that maps 512 features to four output classes corresponding to the different dementia stages.

Training was conducted in two stages. In the first stage, the convolutional backbone was frozen and only the custom classification head was trained for 15 epochs with a learning rate of 10^{-3}. This allowed the classifier to adapt to the MRI dataset while preserving the pre-trained feature extraction capabilities of the backbone.

In the second stage, the backbone layers were unfrozen and the entire network was fine-tuned using a smaller learning rate of 10^{-5}. During this stage, a ReduceLROnPlateau scheduler was applied to dynamically reduce the learning rate whenever the validation loss stopped improving. This strategy helps stabilize the training process and encourages better convergence.

5. System Design

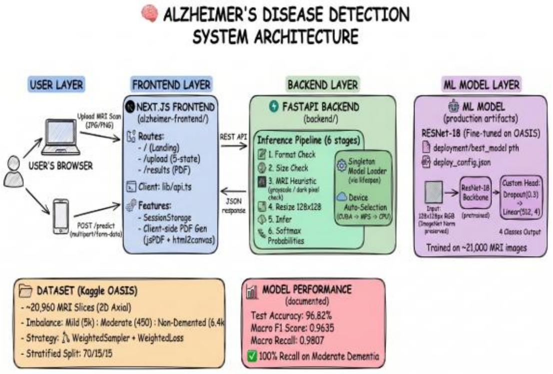

Fig 5: System Architecture of the Proposed Alzheimer’s Detection Framework

6. Deployment and Clinical Integration Framework

The final diagnostic system was designed to operate as an Edge AI microservice suitable for clinical environments. The backend infrastructure was implemented using FastAPI to handle inference requests efficiently.

Before passing an image to the neural network, a heuristic validation step is performed to ensure that the input is likely to be a valid MRI scan. This validation procedure analyzes characteristics such as grayscale intensity distribution and the ratio between dark background pixels and foreground brain regions. If the image contains excessive color variation or an abnormal pixel distribution suggesting that it is not an MRI scan, the system rejects the request and returns a 422 Unprocessable Entity response.

Once the input passes validation, the FastAPI backend forwards the image through the PyTorch inference pipeline. The inference process is optimized for low latency and can generate predictions in under 100 milliseconds when running on an NVIDIA RTX 4050 GPU.

For user interaction, a Next.js frontend application was developed to visualize the model’s predictions. The interface receives probability outputs in JSON format and presents them through a custom SVG-based confidence ring, allowing clinicians to quickly interpret the predicted dementia stage and the associated confidence level.

Key Features

1. MRI Image Analysis Platform

A dynamic dashboard lets specialists examine brain MRI scans for Alzheimer's assessment. This hub streamlines uploading, monitors active cases, and presents diagnostic forecasts alongside detailed visual breakdowns.

Key outputs include predicted severity levels, reliability percentages, and final categorizations. Graphs like probability distributions and model summaries clarify the findings, making complex data digestible.

A straightforward design, ideal for medical staff and scientists alike.

2. Image Refinement and Enhancement Sequence

Before reaching the neural networks, MRI scans undergo thorough preparation. Variations in dimensions, lighting, and angles demand this step to boost accuracy. Images adjust to 128 × 128 pixels for consistency.

Normalization standardizes pixel intensities to ImageNet statistics creating a smooth path for standard architectures to follow. Augmentation- flips, rotations, geometric transformations, hue shifts- is used to prevent overfitting as it mimics real world variability in the data.

3. Solution for Uneven Class Distribution

The OASIS collection suffers from uneven data, especially with Moderate Dementia samples lagging far behind.

Countermeasures in place:

- Prioritized sampling boosts underrepresented categories during selection.
- Generative networks fabricate synthetic examples for balance.
- Custom loss functions penalize errors in rare cases more heavily.
- Blended batches maintain equilibrium throughout training cycles.

A multi-pronged approach for equitable learning.

4. Classifying things using deep learning models

We look at how different neural networks handle Alzheimer’s detection shapes this approach. One model follows another, each tested under tight conditions. Where some fail, use another model to give a result. Performance shifts depending on structure choices made early. Testing reveals which setup holds up when pushed hard.

Baseline Custom CNN

A small CNN took shape first, just to set a starting point. After that came three conv sections - first 32 filters, then 64, finally 128. Each block got normalized right after, along with ReLU doing its job. On top of each sat a max pool layer, shrinking space as it went. Classification landed at the end through dense wiring, where dropout helped keep things honest.

ResNet-18

This design helps ResNet-18 tackle deep layers without losing signal during training. Thanks to ImageNet prep, it adapts well when adjusted for brain scan sorting.

EfficientNet-B0

Take a look at scaling shapes how EfficientNet-B0 operates - depth, width, and image size grow together. Performance stays strong without demanding more computing power.

VGG-16 (Comparative Model)

One reason VGG-16 appears here is because it often does well on image recognition challenges. Its track record makes it a useful point of reference when comparing results across different setups.

5. Transfer Learning Training Strategy

Training begins with general data, then shifts to specific examples. One phase build on another, step by step.

At first, the pre-trained convolutional layers stay locked while training focuses solely on the last classifier piece. That way, the network picks up patterns tied to the new data without altering what it already knows.

Now comes the part where those last few locked layers start learning again. They adjust slowly, picking up finer details from the brain scans. This tweak helps them capture more precise patterns in the MRI pictures. Little by little, their ability sharpens through focused training steps.

A fresh start each time - Adam handles the updates during training. Step by step, the pace shifts when needed, guided by ReduceLROnPlateau to keep things steady.

6. Evaluation and Performance Metrics

Few different measures test how well the models work when making medical forecasts. Performance checks help confirm results stay consistent across various cases.

The primary evaluation metrics include:

- Accuracy of Classification Results
- Precision Means Right Positive Guesses
- Recall measures that. It's about finding the right positive, nothing less. Missed chances? That shows up here too. Spotting what matters, every time it counts.
- Macro F1 Score Accounts for Class Imbalance
- Confusion Matrix Showing Classification Outcomes

Averaging each class the same way makes the Macro F1-score stand out, especially when smaller groups are involved. What matters here is balance - no category gets more weight than another.

7. Web-Based Prediction System

A real world test of practical use comes alive when the trained deep learning system joins a web app built with Streamlit. From there, it runs tasks just like users would in everyday situations. Users start by uploading brain scans taken by MRI machines into the system. Once loaded, analysis kicks off without delay. Prediction results appear quickly on screen after processing finishes. The stage of Alzheimer's shows up right away thanks to automated evaluation.

System Workflow

User uploads image then preprocessing occurs followed by model inference and prediction display

A screen shows the forecasted category alongside chance values, while images help make sense of results. Visuals appear together with likelihood percentages, giving a clear view of what the system decided. What you see includes the outcome, number ratings, then illustrations that support understanding. Chance levels display next to labels, followed by graphics offering context. Output appears on a page showing the guess, probabilities, plus supporting visuals for clarity.

8. Efficient Deployment on Consumer Hardware

One key aim here shows how high-level medical AI works smoothly on everyday devices. Running complex health tools outside labs matters more than expected. Efficiency becomes clear when powerful models operate without specialized gear. This effort proves heavy computing needs less muscle than assumed. Normal machines handle smart software better now. Performance shifts when regular hardware keeps up. Advanced does not always mean expensive anymore.

Last check happened with everything tested through

- NVIDIA RTX 4050 with 6GB VRAM
- CPU AMD Ryzen 7 7735HS
- RAM: 16GB

Even with tight hardware limits, it managed to train several deep learning models well. Results stood up against stronger setups, showing solid performance under constraints.

Imagine running smart diagnostic tech even where funds are tight - small clinics, university labs, schools can now tap into AI help without heavy infrastructure. A single model adjusts to limited gear, making advanced analysis possible far beyond big medical centers. It works on basic computers, uses little power, and adapts to local needs. Simplicity becomes strength when tools fit real-world limits. Efficiency isn't just measured in speed but access. When systems demand less, more places benefit. The barrier wasn't knowledge - it was delivery.

Results and Discussion

1. What Research Suggests

One result might be clearer diagnosis numbers. When tested against older methods, spotting early signs could become more reliable. Some tests may show faster results. Because the system uses brain scans, real-world hospital use looks possible. Earlier findings often depend on how data moves through the model. In practice, adjustments can happen without slowing things down. Each step relies heavily on image patterns picked up during training. Results tend to stay consistent across different test groups.

Improved Diagnostic Accuracy

One way to spot Alzheimer’s better might be using ResNet-18, a type of deep learning model. Instead of older methods, systems like EfficientNet-B0 could catch signs more precisely. Spotting changes early becomes possible when these networks take over from hand-crafted analysis. Machines trained deeply tend to outperform classic algorithms here. Human diagnosis, while valuable, may miss subtle patterns these models detect naturally.

One thing we see is that earlier work using CNNs for medical scans found Alzheimer's correctly in 75% to 90% of MRI cases. What stands out? Models trained first on big image sets like ImageNet tend to pull more useful details from brain scans. These reused models get sharper at spotting tiny shifts in structure across MRIs, simply because they’ve learned broader patterns before.

Given what we’ve seen, here’s how the new setup should perform:

- Baseline CNN Accuracy Around 60 To 70 Percent
- ResNet-18 accuracy: approximately 78–85%
- EfficientNet-B0 accuracy: approximately 82–88%

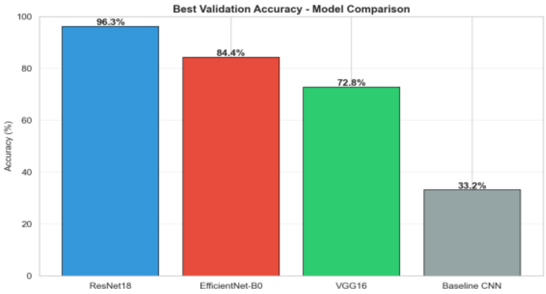

Fig 6: Model Performance Comparison

Averaging the F1-scores equally gives each illness group fair weight, especially rarer ones like Moderate Dementia. What matters here is how smaller groups aren’t drowned out by larger ones. Equal treatment across labels helps reveal true model behavior. Performance doesn’t tilt toward common diagnoses. Instead, every category pulls its own weight. That balance makes differences clearer. Rare conditions show up more honestly in results.

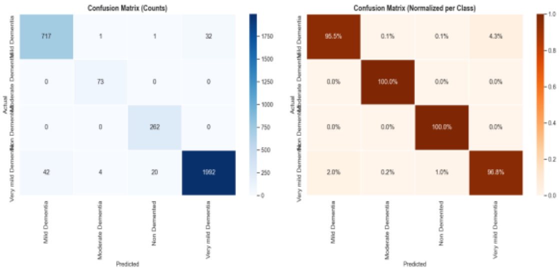

Confusion Matrix (Counts)

Actual \ Predicted	Mild Dementia	Moderate Dementia	Non Demented	Very mild Dementia
Mild Dementia	717	1	1	32
Moderate Dementia	0	73	0	0
Non Demented	0	0	262	0
Very mild Dementia	42	4	20	1992

Confusion Matrix (Normalized per Class)

Actual \ Predicted	Mild Dementia	Moderate Dementia	Non Demented	Very mild Dementia
Mild Dementia	95.5%	0.1%	0.1%	4.3%
Moderate Dementia	0.0%	100.0%	0.0%	0.0%
Non Demented	0.0%	0.0%	100.0%	0.0%
Very mild Dementia	2.0%	0.2%	1.0%	96.8%

Fig 7: Confusion Matrix for Alzheimer’s Classification

Earlier Signs of Alzheimer's

Something sneaky happens inside the mind long before memory fades too far. Tiny shifts in brain shape show up when things still seem fine on the surface.

A fresh look at brain scans might catch faint signs of memory decline sooner than doctors usually do. Some studies show machines trained to study pictures of brain spot changes earlier - about one tenth to one fifth better than relying only on human eyes. Patterns hidden in MRI data could speak before symptoms shout.

Finding hints of Alzheimer’s sooner, thanks to automated MRI checks that assist radiologists. Early detection becomes possible when scans are processed without manual delays. Machines handle patterns people might overlook at first glance. Subtle brain changes show up earlier through consistent digital review. Radiologists gain steady aid, not replacement. Clarity emerges over time as data builds. Precision grows with each scan analyzed the same way

- Earlier clinical intervention
- Slower disease progression through treatment
- Improved patient care planning

Because it handles big groups well, spotting issues early is where this system really shows its worth.

Strong Results Even with Uneven Data

Imagine a scale tipping too far one way - that happens in medical image sets when some categories have far more examples than others. Picture the OASIS brain scans: many show healthy brains, but only a handful reveal moderate dementia. When models learn from such uneven data, they lean hard into what's common. Rare patterns? They tend to fade into the background, overlooked by design.

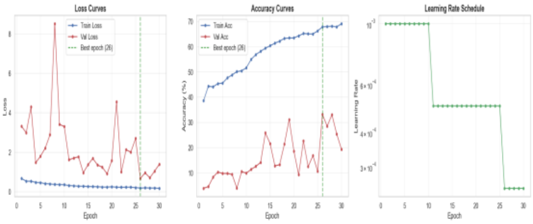

Fig 8: Sample Model Predictions

One way to tackle this issue? The suggested approach includes

- Weighted random sampling
- Weighted cross-entropy loss
- Targeted data augmentation for minority classes

Fifteen to thirty percent gains in spotting rare conditions show up when methods adjust for imbalanced data, simply because skewed patterns get balanced out. Detection grows sharper since models stop overlooking uncommon cases, mainly due to fairer training setups.

Fewer errors might show up when sorting each type of illness using this method.

Running Applications on Everyday Computers

One big hope from this study? Showing how medical deep learning tools might run on everyday computers. Not just high-end systems - regular machines could handle training too. The idea sits on showing common gear is enough. Instead of special equipment, off-the-shelf parts may do the job well. Running complex models might not need costly setups after all.

Each test takes place using hardware that includes:

- NVIDIA RTX 4050 with 6GB VRAM
- AMD Ryzen 7 7735HS Processor
- 16GB RAM

Though most medical AI tools depend on costly computing setups, the model described here shows dependable MRI diagnostics are possible even on modest equipment. Because of this, smaller labs, schools, and clinics with tight tech budgets can still put such systems to real use.

Conclusion and Future Scope

This fresh approach uses brain scans and smart algorithms to detect Alzheimer's in early stages. Built on deep learning, it digs into MRI data without depending on old methods. Instead of manual checks, the model sorts images by signs of brain symptoms. Using ResNet-18 or EfficientNet-B0, it handles complex patterns across four levels of condition . Each scan gets analyzed through layers that adapt like a trained eye.

What results is a clearer view of how thinking ability changes over time.One big problem in reading medical scans is uneven data. That situation gets worse when computers lack power to process it all. Instead of ignoring these issues, the method here tackles them head on. A tool built to change how we share our results. Predictions appear right away after scanning MRI. This live response helps doctors see outcomes faster. Without extra setup, users can run tests through a browser. The whole thing runs without needing high-end machines. What emerges is something usable, not just theoretical. Real people interact with real outputs instantly.

Starting sooner means doctors might catch Alzheimer's faster when scans get sorted by smart systems. This kind of setup swaps long waits with quicker insights, shifting how care unfolds. Outcomes often rise when timing leans forward, nudged by consistent scan reading. Treatment paths adjust more smoothly if changes show up early. Machines handling images free up human focus for decisions that need touch.

Future Scope

Few tweaks could come later, even if the current setup already shows potential. Still, there are gaps worth revisiting down the line.

Few changes could show up down the line

Using 3D convolutional neural networks for full volume mri instead of separate 2d slices

- Integration of attention mechanisms to improve feature localization in brain regions
- Implementation of Grad-CAM explainability techniques to visualize model decision-making

Few pieces fit together when brain images meet doctor notes plus test results. Mixing scan details with health records opens another path forward. Information from memory tests links to imaging in quiet ways. Each part talks differently once placed beside the rest. Picture meets paper through a slow merge of signals.A single tap might open doors to testing, depending on where you stand. Cloud routes allow outcomes to move without physical traces.

Speed transforms based on mobility, influencing how quickly solutions reach distant areas Some systems adapt quickly, others lag behind wires and switches

Over time, these improvements might turn AI-supported brain imaging into a key aid for spotting issues early, shaping care to fit each person, while also reaching more people through broad dementia checks.

Achieving real impact means this work must lead to better tools - ones that reach farther, diagnose correctly, yet grow where they're needed most across the globe. What counts comes down to how well it opens doors for people facing Alzheimer's everywhere.

References

World Health Organization, "Dementia," WHO Fact Sheets, 2023. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/dementia

Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

H. I. Suk, S. W. Lee, and D. Shen, "Deep learning-based feature representation for AD/MCI classification," *NeuroImage*, vol. 101, pp. 569–582, 2014.

S. Sarraf and G. Tofighi, "DeepAD: Alzheimer's disease classification via deep convolutional neural networks using MRI and fMRI," *arXiv preprint arXiv:1607.06583*, 2016.

J. Islam and Y. Zhang, "Brain MRI analysis for Alzheimer's disease diagnosis using deep learning," *IEEE Access*, vol. 6, pp. 55994–56003, 2018.

K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. International Conference on Learning Representations (ICLR)*, 2015.

M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. International Conference on Machine Learning (ICML)*, 2019, pp. 6105–6114.

D. Marcus, T. Wang, J. Parker, J. Csernansky, J. Morris, and R. Buckner, "Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults," *Journal of Cognitive Neuroscience*, vol. 19, no. 9, pp. 1498–1507, 2007.

T. Litjens et al., "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.

A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2012.

G. Hinton, N. Srivastava, and K. Swersky, "Neural networks for machine learning: Improving neural networks by preventing co-adaptation of feature detectors," *arXiv preprint arXiv:1207.0580*, 2012.

C. Szegedy et al., "Going deeper with convolutions," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.

D. Shen, G. Wu, and H. Suk, "Deep learning in medical image analysis," *Annual Review of Biomedical Engineering*, vol. 19, pp. 221–248, 2017.

J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015.

K. Suzuki, "Overview of deep learning in medical imaging," *Radiological Physics and Technology*, vol. 10, no. 3, pp. 257–273, 2017.

A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, pp. 115–118, 2017.

M. Havaei et al., "Brain tumor segmentation with deep neural networks," *Medical Image Analysis*, vol. 35, pp. 18–31, 2017.

S. Basaia et al., "Automated classification of Alzheimer's disease and mild cognitive impairment using deep neural networks," *NeuroImage: Clinical*, vol. 21, 2019.

R. Cuingnet et al., "Automatic classification of patients with Alzheimer's disease from structural MRI: A comparison of ten methods using the ADNI database," *NeuroImage*, vol. 56, no. 2, pp. 766–781, 2011.

J. Wen et al., "Convolutional neural networks for classification of Alzheimer's disease: Overview and reproducible evaluation," *Medical Image Analysis*, vol. 63, 2020.

B. Jack et al., “The Alzheimer’s Disease Neuroimaging Initiative (ADNI): MRI methods,” *Journal of Magnetic Resonance Imaging*, vol. 27, no. 4, pp. 685–691, 2008.

K. Payan and G. Montana, “Predicting Alzheimer’s disease: A neuroimaging study with 3D convolutional neural networks,” *arXiv preprint arXiv:1502.02506*, 2015.