International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Explainable AI-Powered Dementia Detection and Clinical Decision Support Platform

Dr. Sangram S. Dandge
Department of Computer Science and Engineering
Prof. Ram Meghe Institute of Technology and Research
Badnera, India

Harshita Nilesh Rathod, Akshit Sharma, Aditya Sagotra, Shravani Tanaji Pawar
Department of Computer Science and Engineering
Prof. Ram Meghe Institute of Technology and Research
Badnera, India

Abstract— Dementia is a progressive condition that gradually affects memory, reasoning and the ability to manage everyday tasks. Its early signs are often subtle, and reading brain MRI scans reliably takes specialist training. In this paper we propose a clinical decision support platform that uses deep learning to sort brain MRI scans into four dataset-defined stages: Non-Demented, Very Mild, Mild and Moderate. A dementia percentage module converts the model's softmax output into a stage-wise percentage, an overall dementia percentage, a mild-dementia percentage and an early-stage percentage, with a configurable risk flag. Grad-CAM heatmaps show which parts of a scan influenced each prediction; they help inspect model behaviour but do not by themselves establish a diagnosis. A retrieval-augmented generation (RAG) assistant answers medical questions from indexed literature, and the platform builds reporting, specialist search, appointment booking and follow-up tracking around the prediction. OASIS and ADNI are the datasets under consideration, and the system is planned around React, FastAPI, PyTorch and PostgreSQL. At this stage the work is a system design: no model has been trained or evaluated, and no accuracy results are reported.

Keywords— Dementia detection, brain MRI, deep learning, explainable AI, Grad-CAM, Alzheimer's disease, retrieval-augmented generation, medical AI, clinical decision support, dementia percentage, healthcare informatics.

I. INTRODUCTION

A. BACKGROUND

Dementia is a progressive neurological syndrome in which memory, reasoning, language and decision-making gradually decline, to the point where daily activities become difficult. The World Health Organization estimates that 57 million people were living with dementia in 2021, with nearly 10 million new cases each year; Alzheimer's disease is the most common form and accounts for an estimated 60–70% of cases [1]. Early symptoms are easy to overlook, so diagnosis is often delayed. MRI shows the structure of the brain and can reveal the atrophy patterns linked to neurodegenerative disease, but interpreting it needs specialised expertise and takes time.

Open datasets such as the Open Access Series of Imaging Studies (OASIS) [2] and the Alzheimer's Disease Neuroimaging Initiative (ADNI) [3] make it possible to develop and test computational methods on real MRI data, and Convolutional Neural Networks (CNNs) have been shown to learn useful patterns directly from such images for Alzheimer's classification [4], [5]. A class label on its own, however, gives a clinician little to check it against. Grad-CAM [6] marks the image regions that influenced a prediction, and activation-based explanations have also been applied to 3D Alzheimer's classifiers [7]. Retrieval-Augmented Generation (RAG) [8] lets a language model base its answers on retrieved documents.

B. MOTIVATION AND PROBLEM STATEMENT

A diagnosis is only one step in a patient's care. The patient still has to understand the result, find reliable information, reach a suitable specialist and keep up with follow-up visits, often across several unconnected tools and at a stressful time. Most existing work stops at classifying the disease, so even an accurate model does not by itself lead to timely care. A bare stage label also hides how strongly the model supports its decision. We therefore see a need for a single system that takes the user from a prediction, with its supporting probabilities, to the practical next steps.

C. OBJECTIVES

The objectives of this work are to:

- combine prediction, explanation, knowledge retrieval, specialist search and follow-up care in one platform instead of a standalone classifier;
- show a Grad-CAM heatmap and class probabilities beside every prediction, with a clear statement of what they do and do not show (Section IV-E);
- quantify the model output as stage-wise, overall dementia, mild-dementia and early-stage percentages with explicit formulas (Section IV-D);
- answer patient and doctor questions from trusted medical sources through a RAG assistant; and help users find nearby neurologists, book appointments and track.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

D. CONTRIBUTIONS AND PAPER ORGANISATION

Because the work is at the design stage, its contributions are design-level: (i) an integrated twelve-stage workflow and data model that links MRI-based stage classification with explanation, retrieval, reporting and care navigation; (ii) a dementia percentage module that turns softmax probabilities into clinically readable percentages and a configurable risk flag; (iii) a documented rationale for each technology in the stack (Table IV); and (iv) an evaluation protocol fixed before any experiment is run, so that later results can be reported without bias. The system is meant to support clinicians, not to replace clinical diagnosis.

Section II reviews related work, Section III states the requirements, Section IV presents the proposed system, Section V gives the evaluation protocol and project status, Section VI discusses benefits, limitations and applications, and Section VII concludes.

II. RELATED WORK

A. PUBLIC NEUROIMAGING DATASETS

Marcus et al. released OASIS, a public set of cross-sectional MRI scans from young, middle-aged, non-demented and demented older adults [2]. It is easy to access for research, though it may not represent every clinical population. Weiner et al. describe ADNI, which combines longitudinal MRI and PET imaging with CSF and clinical and cognitive measurements [3]; it is a rich multimodal resource, but it was built for research and is not a clinical application. Our four stage labels correspond to Clinical Dementia Rating levels as we understand the OASIS documentation, whereas ADNI groups subjects as cognitively normal, mild cognitive impairment or Alzheimer's disease; if both datasets are used, the labels must be mapped and this mapping confirmed from the dataset documentation.

B. DEEP LEARNING FOR MRI-BASED CLASSIFICATION

Nawaz et al. applied a two-dimensional deep CNN to ADNI MRI data and classified three groups (Alzheimer's disease, mild cognitive impairment and normal control), so their label set differs from our four stages and their work covers classification only [4]. Khagi and Kwon studied how CNN behaviour changes when moving from 2D to 3D designs for MRI and PET, reporting memory use, parameter count and running time [5]. We note that 3D volumes are generally more demanding to process than 2D slices; this is our own reasoning rather than a finding of [5].

C. EXPLAINABILITY IN NEUROIMAGING MODELS

Selvaraju et al. proposed Grad-CAM, which uses gradients flowing into a convolutional layer to localise the image regions behind a CNN's prediction [6]. Yang et al. compared three ways of explaining deep 3D CNN predictions for Alzheimer's classification and observed that the activation-based maps are limited by the resolution of the convolutional layer [7]. Neither work addresses clinical knowledge or patient management.

D. RETRIEVAL-AUGMENTED GENERATION

Lewis et al. introduced RAG, in which a generator is conditioned on documents returned by a retriever for knowledge-intensive tasks [8]. Their system trains the retriever and generator together on general-domain text. Our assistant will not train either component; it will use pretrained embeddings, vector search and a language model, so it is inspired by [8] rather than a reproduction of it, and it needs domain-specific content (research papers, guidelines and educational material) to be useful in medicine.

E. RESEARCH

As Table I shows, each reviewed study addresses one component: a dataset, a classifier or an explanation method. What is still missing is a platform that connects prediction, explanation, evidence-based retrieval, reporting and care navigation. Our design is aimed at this gap; whether it succeeds in practice has not yet been tested.

TABLE I. SUMMARY OF REVIEWED WORKS AND IDENTIFIED GAPS

Ref.	Work	Method / contribution	Identified gap
[1]	Marcus et al., OASIS	Public cross-sectional MRI dataset of non-demented and demented adults	May not represent every clinical population
[2]	Weiner et al., ADNI	Longitudinal MRI, PET, CSF and clinical/cognitive measures	Research resource, not a clinical application
[3]	Nawaz et al.	2D deep CNN on ADNI MRI; three classes (AD, MCI, NC)	Classification only; different label set from ours
[4]	Khagi and Kwon	2D-to-3D CNN study for MRI/PET; reports memory, parameters, run time	No clinical workflow; 3D cost is a design consideration
[5]	Selvaraju et al.	Grad-CAM: gradient-based localisation for CNN predictions	Coarse maps; gives no clinical knowledge or patient management
[6]	Yang et al.	Three visual-explanation methods for 3D CNN Alzheimer's classifiers	Explanation only; maps limited by layer resolution
[7]	Lewis et al.	RAG: retriever plus generator for knowledge-intensive tasks	General-domain; needs medical content

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

III. PROBLEM ANALYSIS AND REQUIREMENTS

A. FUNCTIONAL SCOPE

The platform should classify brain MRI scans into four dataset-defined labels, explain each prediction visually, report class-wise percentages, answer questions from indexed medical documents, draft clinical reports for a clinician to review, manage patient records and MRI viewing, and support specialist search, appointments, follow-up and multilingual assistance. There are two user roles, patient and doctor, plus patient and administrative dashboards. The labels and predictions are research outputs and not a standalone clinical diagnosis.

B. ENTITIES, ATTRIBUTES AND FUNCTIONALITY

We identified eight main entities (Table II). They form the basis of the structured data stored in PostgreSQL.

TABLE II. PRINCIPAL ENTITIES AND ATTRIBUTES

Entity	Principal attributes
Patient	Patient ID, demographics, medical and family history, lifestyle, cognitive assessment history, medications
MRI Scan	Scan ID, patient ID, image/DICOM information, preprocessing status
AI Prediction	Predicted stage, probability distribution, confidence, risk level, model version, heatmap
Doctor	Doctor ID, name, qualification, speciality, experience, fee, hospital, languages, availability
Hospital / Healthcare Centre	Name, location, speciality, distance, contact information, operating hours, services
Appointment	Patient, doctor, date, time slot, status, confirmation, cancellation/rescheduling details
Clinical Document	Document ID, title, source, type, extracted text, embeddings
Clinical Report	Patient information, MRI summary, prediction, confidence, stage, heatmap, observations, recommended follow-up

C. OUTPUTS GENERATED

The outputs follow the workflow stages in Section IV-A: preprocessing status; stage probabilities with a confidence score, dementia percentages and risk level; a Grad-CAM heatmap; a structured clinical report; answers shown with their retrieved source passages; doctor and hospital results; appointment and follow-up information; and patient and administrative dashboards.

IV. PROPOSED SYSTEM

A. OVERVIEW AND WORKFLOW

The platform follows a workflow of twelve stages (Fig. 1, Table III): stages 1–4 cover access and data acquisition, stages 5–7 the AI analysis, and stages 8–12 clinical output and care navigation. A patient or doctor logs in and the patient is registered with a medical history; an MRI scan is uploaded and passes image validation and preprocessing. A deep-learning model then assigns one of the four stages with a confidence score, and Grad-CAM highlights the contributing regions. The platform drafts a clinical report for the clinician to review, the RAG assistant answers questions, and users can search for doctors and hospitals and book appointments with follow-up management. Everything is gathered in a patient or doctor dashboard.

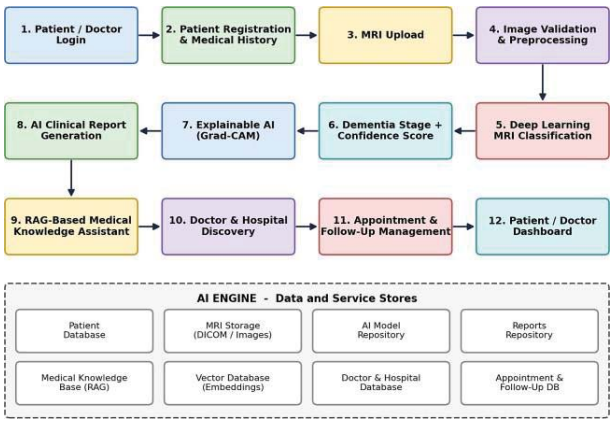


Fig. 1. Proposed twelve-stage AI-assisted healthcare workflow and supporting data/service stores.

TABLE III. THE TWELVE WORKFLOW STAGES

No.	Stage	Description
1	Patient / doctor login	Users authenticate before accessing the platform
2	Registration and medical history	Demographic, medical, family and lifestyle history recorded
3	MRI upload	Brain MRI scan uploaded and stored (DICOM or image)
4	Validation and preprocessing	Scan validated and prepared for the model
5	MRI classification	Trained model produces one score per stage
6	Stage, confidence, percentages	Probabilities, confidence and dementia percentages computed
7	Explainable AI	Grad-CAM highlights regions behind the prediction
8	Clinical report	Structured draft report for clinician review
9	RAG assistant	Questions answered from indexed medical documents
10	Doctor and hospital discovery	Ranked by speciality, location, availability
11	Appointments and follow-up	Booking, medication and reassessment tracking
12	Dashboards	Patient and doctor views of results and appointments

B. DATA AND SERVICE STORES

The AI engine works with the eight stores shown in Fig. 1. We separate them by data type: large image files (DICOM and exported images) and trained model files are kept as files, because a relational database is a poor place for them; structured records such as the entities in Table II go in PostgreSQL, whose transactions help prevent two patients.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

B. TECHNICAL ARCHITECTURE

The architecture has five loosely coupled layers (Fig. 2): a React and TypeScript frontend, a FastAPI backend with JWT authentication, a PyTorch AI layer, a LangChain-based RAG layer using FAISS or ChromaDB, and a PostgreSQL database. The layers are separate so that the AI and RAG components can be developed, evaluated and replaced independently of the web application. Table IV records why each technology is included.

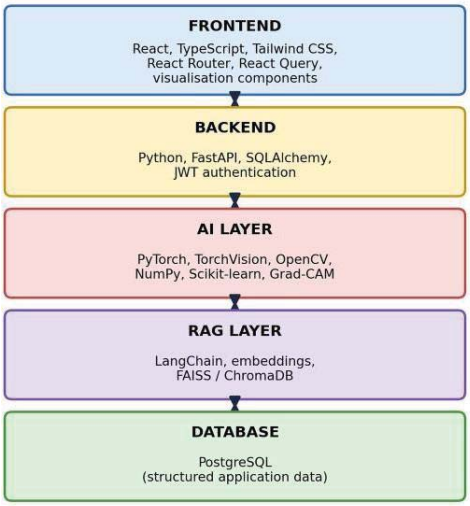


Fig. 2. Layered technical architecture of the proposed platform.

TABLE IV. TECHNOLOGY CHOICES AND RATIONALE

Technology (role)	Why it is included
PyTorch, TorchVision (training, pretrained backbones) [9]	Flexible, gives direct access to gradients (needed for Grad-CAM) and provides pretrained models
OpenCV, NumPy (image handling)	Resizing, intensity normalisation and heatmap overlay
Scikit-learn (metrics, splits) [10]	Tested implementations of precision, recall, F1, ROC-AUC and confusion matrix
LangChain (RAG orchestration)	Connects chunking, retrieval and prompting; adds a dependency whose interfaces change often
FAISS or ChromaDB (vector search)	FAISS is a fast similarity-search library; ChromaDB adds persistence and metadata filtering. The choice awaits a prototype
FastAPI, SQLAlchemy (backend)	Python matches the ML code; request validation and API docs come built in; SQLAlchemy maps Table II to PostgreSQL
JWT (authentication)	Stateless tokens for patient/doctor roles; must be short-lived and sent over HTTPS
PostgreSQL (structured data)	Relational integrity and transactions, e.g. for appointment slots
React, TypeScript, Tailwind CSS, React Query (frontend)	Components suit dashboards and the MRI viewer; types catch API mismatches

C. PREPROCESSING, CLASSIFICATION AND DEMENTIA PERCENTAGES

Every uploaded scan is validated before it reaches the model (Fig. 3). The file must decode as an image; orientation is corrected where applicable; the image is resized to the model input size; intensities are normalised as $x' = (x - \mu)/\sigma$ using the mean and standard deviation of the training data; and the resulting tensor is checked for shape and quality. A scan that fails a check is rejected with a request for a better scan, so unusable images are never classified.

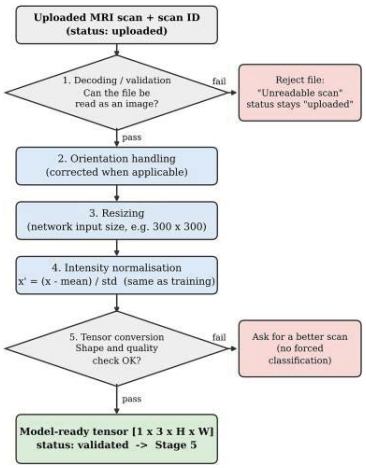


Fig. 3. Workflow of image validation and preprocessing (stage 4).

Each validated scan is assigned to one of four classes (Fig. 4). All candidate backbones are 2D image models, so 3D volumes would first be reduced to slices. EfficientNet-B3
[11] is the primary candidate: it scales depth, width and resolution together for accuracy per parameter. We plan to compare it with DenseNet121 [12], ResNet50 [13] and ConvNeXt [14], and to try a Vision Transformer [15] as an experimental model, since transformers usually need more data or strong pretraining. Because MRI datasets are small, we expect to start from ImageNet-pretrained weights and fine-tune them (transfer learning); natural-image pretraining may not transfer perfectly to MRI, which the comparison will test. The input size is a configurable parameter that will be fixed during tuning. Three data issues must be handled: scans (or slices) from one patient must stay in the same split to avoid leakage; the Moderate class is likely to be rare, so class imbalance needs attention; and stage boundaries such as Very Mild versus Mild are inherently fuzzy.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

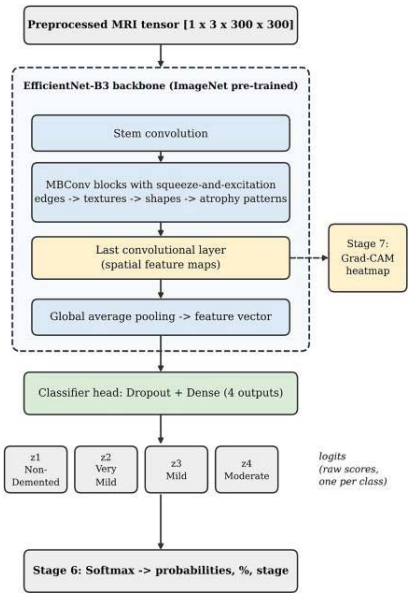


Fig. 4. EfficientNet-B3 classification pipeline (stage 5): backbone, classifier head and four logits.

The classifier outputs one logit z_c for each class $c \in$ {Non-Demented, Very Mild, Mild, Moderate}. Softmax converts the logits into probabilities that are positive and sum to one:

$$p_c = \exp(z_c) / \Sigma_j \exp(z_j) \quad (1)$$

The class percentage, predicted stage and confidence score are

$$P_c(\%) = 100\, p_c;$$
$$\text{Stage} = \text{argmax}_c\, p_c;$$
$$\text{Confidence}(\%) = 100 \max_c p_c \quad (2)$$

The overall dementia percentage is the combined probability of the three dementia stages, which equals 100% minus the Non-Demented percentage:

$$\text{Dementia}(\%) = 100\,(p_{VM} + p_M + p_{Mod}) = 100\,(1 - p_{ND}) \quad (3)$$

The mild-dementia and early-stage percentages are

$$\text{MildDementia}(\%) = 100\, p_M;$$
$$\text{EarlyStage}(\%) = 100\,(p_{VM} + p_M) \quad (4)$$

where ND, VM, M and Mod abbreviate Non-Demented, Very Mild, Mild and Moderate. A case is flagged for priority scheduling when

$$\text{HighRisk} = (\text{Stage} \in \{\text{Mild, Moderate}\})$$
$$\text{OR } (\text{Dementia}(\%) \geq T) \quad (5)$$

where T is a configurable threshold whose value is not yet defined and would need clinical input. The model version is stored with every prediction for traceability.

These percentages have limits that the interface must state. They are model probabilities, not a measure of how much of the brain is affected. Softmax outputs are not guaranteed to be calibrated, so calibration will be checked before the

percentages are shown to users (Section V-B), and rounding is applied only for display.

D. ABILITY MODULE

Explanations are produced with Grad-CAM [6]. For a target class c, the gradient of the class score y^c (taken before the softmax) with respect to the feature maps A^k of a convolutional layer is averaged over the Z spatial positions to give the importance weights

$$\alpha_k^c = (1/Z)\, \Sigma_i \Sigma_j\, \partial y^c / \partial A^k_{ij} \quad (6)$$

The class-specific heatmap is the ReLU of the weighted sum of the feature maps, which keeps only regions that push the score for class c up:

$$L^c = \text{ReLU}(\Sigma_k\, \alpha_k^c\, A^k) \quad (7)$$

The map is upsampled and overlaid on the MRI slice (Fig. 5). It is applied to a late convolutional layer, so it is coarse, a limitation also noted for activation-based maps in [7]; a Vision Transformer would need an adapted variant because it has no convolutional feature maps.

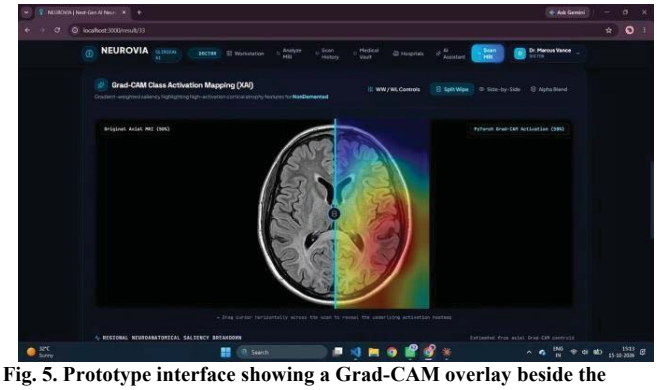


Fig. 5. Prototype interface showing a Grad-CAM overlay beside the original slice. This is an interface screenshot; displayed values are not results from a trained model.

What the heatmap does not show**.** Grad-CAM helps us interpret and debug model behaviour, but it does not prove a diagnosis. It shows where the model's score was sensitive, not why a prediction is correct. A heatmap can fall on scan artefacts, the skull or other non-clinical features, and can look plausible even when the prediction is wrong; explanation methods in health care should not be taken as evidence that a model is reliable [16]. We have not checked whether our heatmaps would agree with clinically known atrophy regions, and doing so would need expert review. The interface will therefore label heatmaps as "model attention, not diagnosis", and any clinical conclusion rests with a qualified clinician.

E. RAG-BASED MEDICAL KNOWLEDGE ASSISTANT

Trusted documents (research papers, clinical guidelines, educational material) are split into passages and stored with title, source, type and embeddings (Table II). retrieval [17] and a similarity-search library [18], and a language model writes an answer from them. Each answer is shown with its source passages so users can check it. The language model has not yet been chosen.

RAG reduces, but does not remove, the risk of unsupported or outdated answers. The assistant will therefore be instructed

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

to answer only from retrieved passages, to say so when nothing relevant is found, and not to give diagnoses. Multilingual assistance is planned, but the method (a multilingual embedding model, or translating queries and answers) is undecided, and translated medical text would need review.

F. REPORT, HEALTHCARE NAVIGATION AND FOLLOW-UP

The clinical report combines the Table II fields for the patient, MRI summary, prediction, percentage block, heatmap and recommended follow-up, and is a draft for clinician review. For care navigation, doctors and hospitals are ranked by speciality, location and availability, so nearby neurologists can be suggested from the AI result; this depends on keeping the directory data accurate, which is not yet planned in detail. Follow-up tracks appointments, medications and reassessments, and cases the risk rule marks as higher risk are flagged for earlier scheduling.

G. EVALUATION METHODOLOGY

Patients will be split into training, validation and test sets, with all scans of one patient kept in one split. The final model will be chosen on validation performance and reported once on the untouched test set, using accuracy, precision, recall, F1-score, ROC-AUC and the confusion matrix. Candidates are EfficientNet-B3, DenseNet121, ResNet50 and ConvNeXt, with a Vision Transformer as an experimental model, trained from ImageNet-pretrained weights on OASIS [2] and, subject to a licence check and label mapping, ADNI [3]. For one class treated as positive against the rest, with TP, TN, FP and FN denoting true positives, true negatives, false positives and false negatives:

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN) \quad (8)$$

$$\text{Precision} = TP / (TP + FP) \quad (9)$$

$$\text{Recall} = TP / (TP + FN) \quad (10)$$

$$F1 = 2 \cdot \text{Precision} \cdot \text{Recall} / (\text{Precision} + \text{Recall}) \quad (11)$$

These are computed per class and then averaged (macro-F1). Overall accuracy can look high when one class dominates, so per-class results matter. ROC-AUC summarises the trade-off between true-positive and false-positive rates, and the confusion matrix shows which stages are mixed up, most importantly neighbouring ones such as Very Mild and Mild.

Because the dementia percentages are softmax probabilities shown to users, we will also check calibration: whether a predicted 80% matches an observed frequency of about 80% correct. A calibration step, for example temperature scaling, will be added if they do not match. Grad-CAM maps and RAG answers will be reviewed qualitatively by domain experts; we do not plan to claim that they are clinically validated.

H. WORKED EXAMPLE OF THE PERCENTAGE MODULE

Table V shows how Eqs. (1)–(4) turn four logits into the reported percentages. The logits are illustrative values chosen to demonstrate the arithmetic. They are not the output of a trained model or of any real patient, and they are not experimental results.

TABLE V. WORKED SOFTMAX EXAMPLE (ILLUSTRATIVE VALUES, NOT MODEL OUTPUT)

Class	Logit z	exp(z)	Prob. p	P (%)
Non-Demented	0.80	2.2255	0.1279	12.79
Very Mild	1.35	3.8574	0.2216	22.16
Mild	2.30	9.9742	0.5730	57.30
Moderate	0.30	1.3499	0.0775	7.75
Total	—	17.4070	1.0000	100.00

The predicted stage is Mild with a confidence of 57.30%. The overall dementia percentage is 22.16 + 57.30 + 7.75 = 87.21%, the mild-dementia percentage is 57.30%, and the early-stage percentage is 22.16 + 57.30 = 79.46% (Fig. 6). The example also shows why confidence and the dementia percentage differ: the model is only 57.30% sure of the Mild stage, yet puts 87.21% of its probability on some form of dementia.

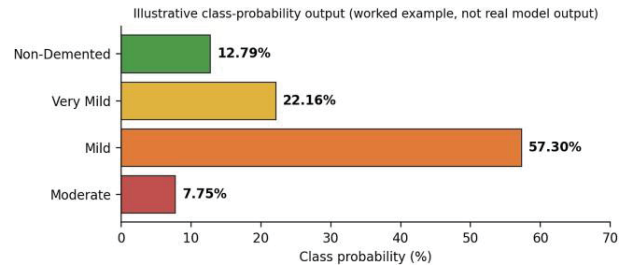


Fig. 6. Class-probability percentages for the worked example (illustrative, not model output).

I. PROJECT STATUS

TABLE VI. COMPLETED VERSUS PLANNED WORK

Status	Item
Completed	Literature review of [2]–[8]; requirements and entity design (Table II); workflow and architecture design (Figs. 1–2); percentage module formulas; prototype interface screens
Planned, no results yet	Dataset selection and licence check; model training and evaluation; calibration and risk threshold T; Grad-CAM implementation and expert review; RAG corpus and assistant; backend and frontend completion; multilingual support; testing

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

V. DISCUSSION

A. EXPECTED BENEFITS

Keeping prediction, explanation, percentages, information search and follow-up in one platform should reduce the effort of moving between separate tools, and showing the full probability distribution gives a clinician more to check than a single label. These are expected benefits of the design and have not been measured.

B. LIMITATIONS AND CONSIDERATIONS

Public datasets may not represent every clinical population [2], and ADNI is a research resource rather than a clinical application [3]. Stage labels come from the dataset, so label noise carries into the model. Slice-based 2D models discard volumetric information that 3D models keep [5]. Grad-CAM explains model behaviour but supplies no clinical knowledge [6] and does not prove a diagnosis [16]; RAG depends on suitable content being indexed [8] and can still produce errors. The dementia percentages are probabilities, not severity, and are meaningful only if the model is calibrated.

Patient records and scans are sensitive, so consent, de-identification, access control and applicable data-protection law must be addressed, and any clinical use would require validation and regulatory approval. The platform is a decision-support tool and not a replacement for clinical diagnosis.

C. APPLICATIONS

Under clinical supervision the platform could serve as a second opinion in neurology clinics, support early-screening review and priority scheduling, guide patients and caregivers through the RAG assistant, coordinate care through specialist search and appointments, and serve research and teaching by comparing architectures and showing how explanation maps are produced.

VI. CONCLUSION AND FUTURE WORK

We have proposed the design of an explainable dementia classification and clinical decision support platform that connects MRI-based stage classification, a dementia percentage module, Grad-CAM explanations, a RAG knowledge assistant and care navigation. The work is currently at the design stage and reports no experimental results. Next, we will evaluate the candidate architectures on appropriately licensed datasets using the protocol in Section IV-H, check calibration and set the risk threshold T with clinical guidance, and then implement and test the remaining modules, including multilingual assistance, with the planned stack.

ACKNOWLEDGMENT

The authors thank their guide, Dr. S. S. Dandge, Dr. M. A. Pund, Head of the Department, and the Department of Computer Science and Engineering, Prof. Ram Meghe Institute of Technology and Research, Badnera, for their guidance and support

REFERENCES

[1] World Health Organization, “Dementia,” WHO Fact Sheet. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/dementia (accessed Oct. 23, 2026).

[2] D. S. Marcus, T. H. Wang, J. Parker, J. G. Csernansky, J. C. Morris, and R. L. Buckner, “Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults,” *J. Cogn. Neurosci.*, vol. 19, no. 9, pp. 1498–1507, Sep. 2007, doi: 10.1162/jocn.2007.19.9.1498.

[3] M. W. Weiner *et al.*, “The Alzheimer's Disease Neuroimaging Initiative: Progress report and future plans,” *Alzheimer's & Dementia*, vol. 6, no. 3, pp. 202–211.e7, May 2010, doi: 10.1016/j.jalz.2010.03.007.

[4] A. Nawaz, S. M. Anwar, R. Liaqat, J. Iqbal, U. Bagci, and M. Majid, “Deep convolutional neural network based classification of Alzheimer's disease using MRI data,” in *Proc. 2020 IEEE 23rd Int. Multitopic Conf. (INMIC)*, 2020, pp. 1–6, doi: 10.1109/INMIC50486.2020.9318172.

[5] B. Khagi and G.-R. Kwon, “3D CNN design for the classification of Alzheimer's disease using brain MRI and PET,” *IEEE Access*, vol. 8, pp. 217830–217847, 2020, doi: 10.1109/ACCESS.2020.3040486.

[6] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 618–626, doi: 10.1109/ICCV.2017.74.

[7] C. Yang, A. Rangarajan, and S. Ranka, “Visual explanations from deep 3D convolutional neural networks for Alzheimer's disease classification,” *AMIA Annu. Symp. Proc.*, vol. 2018, pp. 1571–1580, 2018.

[8] P. Lewis *et al.*, “Retrieval-augmented generation for knowledge-intensive NLP tasks,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 9459–9474.

[9] A. Paszke *et al.*, “PyTorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019, pp. 8024– 8035.

[10] F. Pedregosa *et al.*, “Scikit-learn: Machine learning in Python,” *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.

[11] M. Tan and Q. V. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in *Proc. 36th Int. Conf. Mach. Learn. (ICML)*, PMLR vol. 97, 2019, pp. 6105–6114.

[12] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 2261–2269, doi: 10.1109/CVPR.2017.243.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

[13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 770–778, doi: 10.1109/CVPR.2016.90.

[14] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A ConvNet for the 2020s," in *Proc. IEEE/CVF.*

(This work is licensed under a Creative Commons Attribution 4.0 International License.)