

LIGHTWEIGHT DEEP LEARNING FRAMEWORK FOR DEMENTIA STAGE IDENTIFICATION USING CORONAL MRI SCANS

A Major Project Report
Submitted in Partial Fulfillment of the Requirements for the Degree of

MASTER OF TECHNOLOGY

in

DATA SCIENCE

by

PRATHAM YADAV

24/DSC/22

yadavpratham37@gmail.com

Under the Supervision of

Dr. Divyashikha Sethia

Associate Professor, Department of Software Engineering
Delhi Technological University



Department of Software Engineering

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daultpur, Main Bawana Road, Delhi-110042, India

May, 2026



DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daultpur, Main Bawana Road, Delhi-110042

ACKNOWLEDGEMENT

I would like to express my sincere gratitude to my supervisor, _____

Dr. Divyashikha Sethia

for the continuous guidance, encouragement, and valuable suggestions throughout the course of this thesis work. Her/His insights have been instrumental in shaping the direction and quality of this research.

I also thank the faculty members of the Department of Software Engineering for their academic support and encouragement. Special thanks to the members of the Data Science program for creating a collaborative research environment.

I am grateful to my family and friends for their unwavering support and motivation throughout this challenging project. Their patience and understanding have been invaluable during the demanding periods of implementation, experimentation, and thesis writing.

Place: Delhi

Date: 6/7/2026

Pratham Yadav



DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-110042


CANDIDATE'S DECLARATION

I hereby declare that the report work entitled

Lightweight Deep Learning Framework For
Dementia Stage identification using Coronel MRI
Scans.

submitted in partial fulfillment of the requirements for the award of the degree of Master of Technology in Data Science, in the Department of Software Engineering, Delhi Technological University, is a bona fide record of the research work carried out by me during the period from August 2024 to May 2025 under the supervision of Dr. Divyashikha Sethia.

I further declare that the work presented in this thesis has not been submitted, either in part or in full, for the award of any other degree, diploma, or fellowship of this or any other institution or university.


Candidate's Signature

This is to certify that the candidate has incorporated all the corrections and suggestions recommended by the examiners, and the declaration made by the candidate is true and correct to the best of our knowledge and belief.



Signature of Supervisor(s)

Signature of External Examiner



DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-110042

CERTIFICATE

This is to certify that the Project Dissertation titled "Lightweight Deep Learning Framework for Dementia Stage Identification Using Coronal MRI Scans" submitted by PRATHAM YADAV, Roll No. 24/DSC/22, Department of Software Engineering, Delhi Technological University, Delhi, in partial fulfillment of the requirement for the award of the degree of Master of Technology (Data Science), is a record of the project work carried out by the student under my supervision.

To the best of my knowledge this work has not been submitted in part or full for any Degree or Diploma to this University or elsewhere.

A handwritten signature in blue ink, appearing to read "Divyashikha", with a long horizontal stroke extending to the right.

Dr. Divyashikha Sethia

Associate Professor

Department of Software Engineering

Place: Delhi

Date: May 14, 2026

ABSTRACT

Dementia is a rapidly increasing global health problem and refers to a range of progressive neurodegenerative conditions that result in the loss of memory, cognition and independent functioning. Dementia currently affects more than 55 million people globally and is expected to affect 139 million more by 2050, making the need for efficient, accurate and scalable diagnostic techniques more urgent than ever.

Another important problem is the extreme class imbalance found in typical datasets for dementia imaging, especially the clinical stage Moderate-Demented (MD), which leads to a class-imbalanced diagnostic behavior of models trained on the conventional way.

In this thesis, LW-DMNet (LightWeight Dementia MRI Network) is introduced, a parameter-efficient deep learning network designed for automated multi-class dementia stage classification using coronal MRI images, which is clinically feasible. LW-DMNet stands for a lightweight custom Convolutional Neural Network (CNN) architecture designed from first principles for the features extracted using coronal brain MRI, but using no pre-trained weights. It is built using four hierarchical convolutional blocks, a global average pooling layer, and a three-layer fully connected classification head, with only 13.39 million trainable parameters, which is about 90% fewer than the number of parameters in VGG16 (approximately 138 million).

In order to solve the class imbalance problem in a principled way, the hybrid re-sampling technique SMOTETomek is used. It combines the Synthetic Minority Over-sampling Technique (SMOTE) to create new synthetic samples for the under-represented classes by interpolating them in feature space, with the Tomek Link removal technique to remove samples in ambiguous regions at the class boundary to produce a balanced and well-separated training distribution. The experiment results validate LW-DMNet with an accuracy of 98.44% and F1-scores of 1.00 for the critically important Moderate-Demented class, and 98.19% for macro-averaged scores, despite a small number of 64 training instances for this class. The discriminative power of the proposed model is further validated by the macro-average AUC, which is 0.998. These findings motivate LW-DMNet as a clinically feasible, computationally efficient approach to dementia staging based on brain images.

CONTENTS

Acknowledgement	i
Candidate's Declaration	ii
Certificate	iii
Abstract	iv
List of Figures	3
List of Tables	4
1 Introduction	5
1.1 Deep Learning in Neuroimaging: Promise and Limitations	6
1.2 Problem Statement	7
1.3 Proposed Solution: LW-DMNet	8
1.4 Objectives of the Study	8
2 Technical Background	10
2.1 Dementia: Clinical Presentation and Neuropathological Basis	10
2.2 Magnetic Resonance Imaging in Dementia Diagnosis	11
2.3 Fundamentals of Convolutional Neural Networks	12
2.4 Transfer Learning and Its Limitations in Medical Imaging	13
2.5 Handling Class Imbalance with SMOTETomek	14
3 Related Work	16
3.1 Early Deep Learning Approaches to Dementia Classification	16
3.2 Transfer Learning Approaches	17

3.3	Three-Dimensional and Volumetric Approaches	18
3.4	Lightweight and Task-Specific Approaches	18
3.5	Class Imbalance in Dementia Staging: Literature Overview	19
3.6	Research Gap and Justification for LW-DMNet	19
4	Proposed Methodology and Experimental Results	22
4.1	Dataset Description	24
4.2	Data Preprocessing Pipeline	24
4.3	Data Partitioning and Augmentation	25
4.4	Class Imbalance Handling with SMOTETomek	25
4.5	LW-DMNet Architecture	27
4.5.1	Convolutional Feature Extractor	28
4.5.2	Global Average Pooling	28
4.5.3	Fully Connected Classification Head	29
4.5.4	Parameter Count and Efficiency Analysis	29
4.6	Training Strategy	30
4.7	Experimental Setup and Implementation Details	30
4.8	Overall Test Set Performance	31
4.9	Per-Class Classification Report	31
4.10	Confusion Matrix Analysis	33
4.11	ROC-AUC Analysis	34
4.12	Comparative Analysis with State-of-the-Art	34
5	Discussion and Analysis	37
5.1	Interpretation of Performance Results	37
5.2	Architectural Efficiency and Its Clinical Implications	38
5.3	Significance of SMOTETomek in Addressing Clinical Data Realities .	39
5.4	Limitations of the Present Study	39
6	Conclusion and Future Work	41
6.1	Conclusion	41
6.2	Future Work	42

LIST OF FIGURES

4.1	Refined Flowchart of LW-DMNet Methodology Pipeline	23
4.2	LW-DMNet: Architecture for Dementia Stage Classification from Coronal MRI Images.	27
4.3	The LW-DMNet (Confusion Matrix) on the Held-Out Test Set. Consistently high accuracy classification in all four dementia stages is shown by strong diagonal concentration.	33
4.4	Multi-Class ROC-AUC Curves for LW-DMNet.	34

LIST OF TABLES

3.1	Comparing with Recent Dementia Classification Models.	21
4.1	Distribution of Training Data Before and After SMOTETomek Resampling	26
4.2	Model Complexity Comparison: Trainable Parameters	29
4.3	Overall Test Set Performance Metrics for LW-DMNet	31
4.4	Per-Class Classification Report for LW-DMNet on the Test Set	32
4.5	This is a comprehensive comparison of LW-DMNet with the state-of-the-art dementia. Classification Models.	35

CHAPTER 1

INTRODUCTION

Dementia is not a single disease but rather an overarching clinical syndrome characterized by the progressive deterioration of cognitive functions severe enough to interfere with a person's daily life and activities. It affects many areas such as memory, language, problem-solving, and executive functioning, and is the result of several different underlying neuropathological conditions. The most common form is Alzheimer's disease (AD) which affects around 60–70

The impact of dementia is overwhelming. According to the World Health Organization, 55+ million people are currently suffering from dementia globally [1]. This number is expected to rise almost two-fold to 78 million by 2030 and to 139 million by 2050, with the majority of this increase attributable to demographic ageing in the world. The economic effects are similarly huge; the cost of dementia to the world is estimated to be greater than one trillion US dollars annually, with this cost projected to double over the next 10 years. These costs include direct medical expenses, social support and the unquantifiable informal care costs placed on families and communities.

Early and correct classification of the stage of dementia is of utmost clinical significance in this context. Stage-specific diagnosis allows clinicians to start timely, appropriate pharmacological and non-pharmacological treatments, advise patient and/or family on prognosis, refer patients to relevant clinical trials and plan for long-term treatment. The four known stages of clinical dementia severity (Non-Demented (ND), Very Mild-Demented (VD), Mild-Demented (MiD), and Moderate-Demented (MD)) have different patterns of cognitive and structural brain changes, and making precision distinctions between them is important for personalized patient management.

Structural Magnetic Resonance Imaging (MRI) emerged as the most informative and most widely used neuro-imaging technique for the detection and follow-up of the patterns of brain atrophy characteristic of dementia. T1-weighted MRI can be used to obtain a detailed image of important biomarkers such as global cortical thinning, loss of volume in the hippocampus, and enlargement of ventricles, all of which are associated with the clinical phase of dementia and align well with the MRI's high spatial

resolution and excellent soft-tissue contrast [2]. One of the many planes of acquisition, the coronal plane is especially useful because it provides a direct view of the structures of the medial temporal lobe including the hippocampus and entorhinal cortex, the earliest and most reliably affected structures in Alzheimer’s disease pathology.

Manual interpretation of magnetic resonance imaging (MRI) for dementia staging, however, is a cognitively and complexly demanding task. It takes years of training and there is significant interrater variability, and as the amount of neuroimaging data being produced worldwide continues to increase, it becomes increasingly impractical. This restriction is an urgent need for automated, objective, and scalable computational tools to help clinicians in the diagnosis process, minimize diagnostic variability, and provide expert-level diagnostic capability even in the absence of a specialist radiologist.

1.1 Deep Learning in Neuroimaging: Promise and Limitations

Deep learning (DL) and Convolutional Neural Networks (CNNs) are the most recent paradigm for automated medical image analysis in the past decade. In the medical imaging field, hundreds of applications of deep learning in radiology, pathology, ophthalmology, and other fields demonstrated “remarkable performance” according to a recent survey of Litjens et al. [3]. Dementia classification based on DL has been shown to yield accuracy rates comparable to or superior to other machine learning techniques, and in some studies, comparable to that of expert clinicians for some parts of the classification task in the classification of dementia using neuroimaging.

In the recent literature of DL-based dementia classification, the use of large-scale transfer learning (TL) models are commonly used for dementia classification. These models (VGG-family networks, ResNet variants, InceptionNet, DenseNet and EfficientNet architectures) are pre-trained on a natural image dataset called ImageNet with more than 1.2 million images across 1,000 categories, and then fine-tuned on the target medical imaging dataset. The motivation for this is because the low level feature representations learnt from natural images (edges, textures, gradients) can be transferable to the medical imaging domains and serve as a good initialisation that doesn’t require a lot of labelled medical data.

This has been found to be true for various cases, however, for coronal MRI dementia classification, there are serious drawbacks in using the above architectures. VGG16 is one of the popular transfer learning backbone architectures used for this task, which has approximately 138 million trainable parameters. The number of parameters is approximately 25 million for ResNet50, and approximately 20 million for the more efficient EfficientNet base family. The features of coronal brain MRI image space are

much smaller than the space of natural image classification, where the discriminative signal is found in subtle features such as the volume of the hippocampus, width of the sulci and size of the ventricles.

They are also sensitive to overfitting on the small medical datasets, and need strong regularization, data augmentation and hyperparameter tuning to ensure stable generalization. Further, the latency of inferences is a clinically important concern for large models that are valuable for quickly offering clinical assistance. These factors lead the design of new compact, specialized architectures that can produce diagnostic performance comparable to the original, but with much reduced computational needs.

The second major challenge that is frequently not addressed in the current literature is that of dealing with the imbalanced nature of dementia neuroimaging datasets. The public datasets used are heavily imbalanced, with thousands of samples having labels Non-Demented, and as many as 64 original images having labels Moderate-Demented. A model trained using such imbalanced distributions will learn biased decision boundaries such that it will achieve a higher accuracy on the majority classes while underperforming on the minority classes. Though they show high headline accuracy rates, the clinically significant or end-of-stage care planning step, known as the Moderate-Demented stage, is absent from most models, making them almost useless in practice.

1.2 Problem Statement

The central research question for this thesis is as follows: How to design a deep learning architecture that achieves the state-of-the-art diagnostic accuracy for four classes of dementia from coronal MRI images, with a number of parameters significantly less than current transfer learning approaches, and that performs robustly and fairly across all classes of dementia including the under-represented Moderate-Demented class?

This problem involves three technical challenges which are interrelated. The first is architectural design, that is, the design of a CNN that is specifically tailored to the feature extraction needs of the coronal brain MRI without the need of pre-trained weights or ensemble strategies. The second is data rebalancing where a principled hybrid re-sampling technique not only generates meaningful synthetic data of minority classes but also cleans decision boundary of noisy data. The third is rigorous evaluation, using a complete and unbiased suite of performance measures including per-class precision, recall, F1-score and AUC-ROC analysis, which provide a comprehensive picture of the model's performance across all clinical stages.

1.3 Proposed Solution: LW-DMNet

To address these problems, this work introduces a custom CNN architecture, LW-DMNet (LightWeight Dementia MRI Network), specifically designed for dementia stage classification on a coronal MRI. Constructed from scratch without using pre-trained weights, LW-DMNet is able to learn a small and domain specific feature representation directly from the neuroimaging data. The architecture consists of four consecutive convolutional blocks of differing filter depth (32, 64, 128, and 256) followed by global average pooling and a three-layer fully connected classification head, for a total of 13.39 million parameters that are trainable. It has about 90

To address the problem of class imbalance, LW-DMNet uses the hybrid resampling method SMOTETomek only in the training partition so as to avoid any possible data leakage into the evaluation sets. The SMOTETomek algorithm is a fusion of the SMOTE oversampling algorithm, which generates synthetic minority instances by interpolating between existing minority instances in feature space, and Tomek links cleaning procedure, which identifies and removes borderline instances that lie close to the inter-class boundaries, enhancing the quality and separability of the training distribution.

The experimental results on the Kaggle Alzheimer’s MRI Dataset [4] show the test accuracy of 98.44%, the test F1-score of 98.19% and the macro-average AUC of 0.998. Importantly, the model achieves perfect precision and recall (1.00) for the Moderate-Demented class, indicating the effectiveness of the SMOTETomek rebalancing strategy. The outcomes of this study are in agreement with other works that have shown that LW-DMNet offers a practical and effective alternative to more complex transfer learning-based approaches in automated dementia staging, while avoiding the computational complexity of these methods.

1.4 Objectives of the Study

The specific aims setting the framework of the research in this thesis are listed below. The main goal is to develop and apply a parameter-efficient Convolutional Neural Network (CNN) architecture for the four-class dementia staging problem from coronal T1-weighted MRI images with the absence of pre-trained weight and transfer learning components. The second goal was to tackle the large class imbalance that is common in most dementia neuroimaging data sets by using SMOTETomek hybrid resampling method in the training partition to enhance the sensitivity of the model to the under-represented classes of dementia, specifically the Moderate-Demented class. The third

objective consisted of a thorough analysis of the proposed model with the use of a large set of performance metrics (accuracy, precision, recall, F1-score, AUC-ROC, and confusion matrix analysis) in all four dementia stages to guarantee that the results were unbiased and relevant to clinical use. The fourth aim is to systematically compare the diagnostic performance of LW-DMNet to state-of-the-art competing approaches found in the literature as well as the model complexity with regard to trainable number of model parameters. The fifth objective is to discuss the applicability of the proposed framework to clinical settings, its deployment advantages in low resource settings and directions for future research to create a complete and comprehensible diagnostic support system based on neuroimaging.

CHAPTER 2

TECHNICAL BACKGROUND

This chapter lays the theoretical and technical groundwork on which the proposed will be based. LW-DMNet model is constructed. It starts with a review of the clinical and neuropathological. foundation of dementia, its grading, and the pivotal place of MRI in its diagnosis. It then describes the main concepts of deep learning that the suggested model is based on, such as the design and workings of Convolutional Neural Networks. This is then a discussion of transfer learning, and its use in medical image analysis. and its intrinsic constraints when used in neuroimaging tasks. The chapter concludes including a step-by-step discussion of the imbalance in medical data classes and the theory. of the SMOTETomek hybrid resampling approach used in this work.

2.1 Dementia: Clinical Presentation and Neuropathological Basis

Dementia refers to acquired progressive loss of cognitive impairment severe. to the extent of crippling everyday activities, due to underlying brain disease. It does not occur as a usual age-related effect but occurs as a pathological condition that is based on certain neurodegenerative or cerebrovascular disease mechanisms. The most common etiology is Alzheimer disease which is histopathologically defined by extracellular aggregation of amyloid-beta. plaques and intraneuronal structure of neurofibrillary tangles made of hyperphosphorylated tau protein. These molecular actions cause a cascade of synaptic dysfunction, neuroinflammation, and eventually the loss of neurons, resulting in progressive macroscopic brain. atrophy.

Alzheimer-type dementia is clinical progression that can be broadly classified to four. stages that are commonly used as targets of classification in computational dementia staging studies. The first is Non-Demented (ND) which is associated with cognitively normal individuals where no clinical significant cognitive impairment is involved. Structural MRI of this group would reveal mild age changes and no patterns of pathological atrophy. The second level, Very Mild-Demented (VD), is similar to the clinical concept of Mild Cognitive Impairment (MCI): in this case, patients have subjective and objective experience. fairly severe memory problems out of proportion

to age, but not yet sufficiently serious to be of any effective concern to day-to-day life. The stage is the critical period of therapeutic intervention as it comes before the onset of complete dementia. Structural MRI is initiated by present minor loss of hippocampal and entorhinal cortex in this phase.

The third phase, Mild-Demented (MiD), is marked by a distinctly noticeable cognitive impairment which starts to disrupt everyday capabilities. Memory deteriorates, and language, visuospatial, and executive deficits can occur. emerge. At this phase, MRI is usually characterized by more severe atrophy of the medial temporal lobe, a significant loss of hippocampal volume and a slight loss of cortical thickness. The fourth and the most serious level is Moderate-Demented (MD) which is characterized by severe cognitive impairment that is in need of full-time care. People cannot carry out the simple tasks of daily living and memory loss, disorientation and behavioral abnormalities are the most common. MRI results contain widespread cortical atrophy, considerable ventricular enlargement. and extensive subcortical changes. This capability to automatically and correctly distinguish between these four stages based on MRI data has direct clinical use and is the key task discussed in this thesis.

2.2 Magnetic Resonance Imaging in Dementia Diagnosis

Structural MRI is the new neuroimaging modality of diagnosis and. clinical and research monitoring of dementia [2]. The physics concept behind MRI is the process of measuring nuclear magnetic resonance signatures of hydrogen. protons in tissue water which are manipulated by radiofrequency pulse sequences to produce tissue contrast. T1-weighted sequences are the most popular used in. structural brain studies, give an excellent comparison between gray matter, white matter, and cerebrospinal fluid, and are peculiarly adapted to the detection of the volumetric changes in neurodegenerative disease brain parenchyma.

The coronal acquisition plane is of particular diagnostic relevance in dementia assessment. This plane provides a direct, orthogonal view of the medial temporal lobe structures — specifically the hippocampus, parahippocampal gyrus, and entorhinal cortex — that are known to exhibit the earliest and most consistent atrophy in Alzheimer’s disease. Radiologists and neuroimaging specialists routinely assess hippocampal volume and morphology on coronal slices as a key component of dementia workup, and the visual rating scales commonly used in clinical practice (such as the Scheltens medial temporal lobe atrophy scale) are standardized for coronal orientation. This anatomical and clinical basis justifies the focus on coronal MRI images in the present work.

Quantitative MRI-derived biomarkers of neurodegeneration — particularly hippocampal volume, total brain volume, and cortical thickness measurements — have been validated as sensitive and specific markers for dementia staging. Automated volumetric tools such as FreeSurfer can extract these measures from high-quality T1 MRI scans, but their computational requirements, need for high-quality input data, and the time required for processing make them impractical for routine clinical triage. Deep learning-based approaches offer a complementary strategy that can provide rapid, automated staging directly from raw MRI images without requiring explicit anatomical segmentation.

2.3 Fundamentals of Convolutional Neural Networks

Convolutional Neural Networks are a type of artificial neural network that uses the spatial organization of image data with a sequence of learnable convolutional filters [5]. CNNs use weight sharing and in contrast to fully connected networks, local connectivity, scaling down the number of parameters to process highdimensional image inputs by a very large factor, and also capturing strong inductive biases about the translation-invariant character of visual features.

The basic computing block of a CNN is the convolutional layer, which implements a collection of (learnable) filters (also known as kernels) to the input feature map via a sliding window operation. The dots between the weights of each filter and the local are computed by each filter. response at every spatial location, and outputs a feature map of response that represents the power of the detection of the respective learned feature at every location in the image. The hierarchy in learning can be learnt by piling up several convolutional layers. representations: lower layers encode primitive features, like edges and gradients, intermediate layers represent more complex patterns, e.g. textures and shapes and even deeper layers are high-level semantic representations that are pertinent to the classification problem.

The method of batch normalization is an important regularization/optimization method that is used extensively in current CNNs. The implementation of batch normalization was introduced to solve the internal covariate shift issue in deep networks, normalizing the activations of all the layers in the mini-batch, to be zero-mean and unit-variance, followed by the use of learnable scale and shift parameters. This normalizes the training process, permits higher learning rates, is a mild regularizer, and has proven to consistently enhance both speed of training and end model. When the training data is limited and typically noisy, as is the case in medical image CNNs, it is especially the batch normalization that contributes to the model stability.

Max-pooling (most often referred to as pooling operations) is used to reduce progressively the spatial dimensions of feature maps as they are filtered by successive convolutional blocks. A 2x2 max-pooling layer that uses a stride of 2 helps cut each spatial dimension in half, in effect summarizing the most salient feature activations in local regions. This spatial subsampling allows a form of translation invariance, helps reduce the computational load of the later layers, and has a more modest regularization effect of eliminating accurate spatial information in favor of more high-level feature presence.

The ReLU (Rectified Linear Unit) activation function, which is $f(x) = \max(0, x)$, is the conventional non-linearity of modern CNN models. Its computational simplicity and the lack of the gradient saturation issue of sigmoid and It has been the activation of choice in deep networks due to hyperbolic tangent activations. Randomizing dropout, randomly choosing a certain percentage of neuron activations to zero. during each forward pass in training, provides an effective mechanism for preventing introducing stochastic noise to overfit the network to learn in a stronger way. and overlapping representation of features.

Global Average Pooling (GAP) is an alternative to the traditional practice of flattening the final convolutional feature maps into a fully connected layer. GAP computes the spatial average of each feature map, producing a single scalar per channel and reducing the multi-dimensional feature tensor to a one-dimensional vector regardless of input spatial size. This provides several advantages: it dramatically reduces the number of parameters at the interface between the convolutional backbone and the classification head, it is more robust to spatial translations in the input, and it acts as a structural regularizer that reduces the risk of overfitting in the fully connected component of the network.

2.4 Transfer Learning and Its Limitations in Medical Imaging

Transfer learning refers to the practice of initializing a neural network with weights pre-trained on a large source task (typically ImageNet classification) and then fine-tuning the resulting model on a smaller target task. The theoretical basis for this approach is the hypothesis that the feature representations learned in early and intermediate layers of a deep network are sufficiently general — capturing basic visual primitives such as edges, textures, and color gradients — that they provide useful initial representations for target tasks with different but visually related content.

In the medical imaging domain, transfer learning has been widely adopted as a practical strategy for compensating for the limited availability of annotated training data.

Large annotated medical imaging datasets remain far scarcer than their natural image counterparts due to the cost and expertise required for expert annotation, patient privacy constraints, and the relative homogeneity of medical imaging data compared to the broad visual diversity of natural photographs. Transfer learning from ImageNet pre-trained models has been shown in numerous studies to provide a useful initialization that reduces the labeled data requirement and often improves final model performance compared to training from scratch on small medical datasets.

However, the adoption of transfer learning for medical imaging tasks carries a set of important limitations that are particularly relevant in the context of dementia staging from coronal MRI. The most fundamental concern is over-parameterization. Models such as VGG16, with its approximately 138 million parameters, were designed to distinguish between 1,000 diverse natural image categories with high intra-class visual variability. The dementia staging task, by contrast, requires discrimination between four classes of brain MRI images whose visual differences — primarily in the degree of hippocampal atrophy and cortical thinning — are subtle, localized, and confined to a narrow region of the image. The vast majority of the representational capacity of a VGG16 network is therefore irrelevant to this task and constitutes parameter overhead that increases computational cost, training time, and the risk of overfitting without proportionate diagnostic benefit.

Furthermore, there is a fundamental distributional mismatch between natural RGB photographs and grayscale MRI images. MRI images have a single intensity channel, exhibit a radically different statistical distribution of pixel intensities, and contain domain-specific structural regularity (fixed orientation, consistent anatomical layout) that has no analogue in natural image collections. Converting grayscale MRI images to three-channel format to accommodate ImageNet-pretrained architectures is a common workaround but introduces artificial redundancy and does not resolve the underlying domain mismatch. These considerations collectively motivate the development of compact, task-specific architectures that learn domain-relevant representations from the medical data itself, without inheriting the over-parameterization and domain mismatch of transfer learning baselines.

2.5 Handling Class Imbalance with SMOTETomek

Class imbalance is a widespread issue in medical machine learning, which is caused by the prevalence differences which occur in the patient population between disease states. In the context Standard publicly available datasets have a severe imbalance of dementia staging: the Thousands of instances are normal in Non-Demented class, but

there are the Moderate-Demented. class can contain less than 100. Training a classifier on this imbalanced distribution when it is not corrected appropriately will lead to the classifier being trained to give low probability to the minority. classes not because of any inherent discriminative difficulty but simply because minimizing the training loss on an imbalanced dataset is most easily achieved by predicting the majority class. This contributes to exaggerated total accuracy values which masks weak sensitivity. on the minority classes which are frequently most clinically important.

The SMOTE (Synthetic Minority Over-sampling Technique) algorithm [6] takes care of. class imbalance by creating new synthetic minority-class examples by interpolating in feature space instead of simply duplicating. In each sample of minority classes, SMOTE finds its k nearest classmates, and randomly picks one of these nearest classmates and generates a new synthetic sample at a random point. the line segment between original sample and line neighbor chosen. This interpolation process makes sure that synthetic samples are realistic instances that are part of the feature-space distribution of the minority class, not literal copies, which simply happen to lie in this distribution add number without any new data.

The Tomek Links algorithm [7] offers a complementary undersampling operation. that specifically targets the decision boundary region. A pair of Tomek links is called a pair. of samples of various classes which are nearest neighbors to each other (mutual nearest neighbors) close friend of the other, although of different classes. These pairs of samples are categorized as the borderline or grey cases that are near the inter-class line. and the presence of which may confuse the classifier. By eliminating the member of the majority-class. of each Tomek link, the algorithm clears ambiguous samples from the class boundary, optimally making the resulting training distribution sharper.

A hybrid method called SMOTETomek hybrid method [8] performs both of the operations one after the other: SMOTE. is initially used to oversample the minority group(s) to a desired number, and Tomek Link. then removal is used to eliminate boundary noise of the resulting distribution. This two step process is a solution to the problem of class imbalance on both ends: it boosts the representation. of minor classes and at the same time eliminating the vaguity of samples of majority classes. that the learning process is confused by most. The outcome is that the training dataset is more. balanced classes in a better proportion and more explicitly organized in inter-class division. Importantly, in the current work SMOTETomek is used on the training only. partition; the validation and test sets are kept with their original uneven distribution to. make sure that performance evaluation is based on realistic data distribution of clinical data.

CHAPTER 3

RELATED WORK

Deep learning in automated staging and classification of dementia. based on neuroimaging information has been a dynamic and fast developing field of study since more than a. decade. This chapter will give a detailed overview of the most pertinent and powerful. Research in this field, which locates the suggested LW-DMNet framework into a wider context. research landscape and determining the particular limitations that drive the contributions. of this thesis. The review is presented in thematic format discussing the trend in architectural design, the use of datasets, the problem of the imbalance of classes, and the critical problem of computational. efficiency. A summary comparison of representative prior works is presented in Table 3.1.

3.1 Early Deep Learning Approaches to Dementia Classification

The first attempts at using deep learning to classify MRI-based dementia. The success of CNNs with the general computer vision inspired the need to apply them to data. community after AlexNet sets the record in the ImageNet challenge. in 2012. One of the earliest notable demonstrations of CNN was by Sarraf et al. [5] and repeated neural network-based Alzheimer disease MRI and classification. functional MRI results, which prove the conceptual viability of end-to-end deep learning. for neuroimaging classification. This and other early works illustrated that CNN feature representations, which are learned even with relatively small neuroimaging data. could even obtain diagnostic patterns of brain atrophy.

Litjens et al. [3] gave an in-depth overview of deep learning applications in the field of. medical imaging modalities, reporting the swift transfer of convolutional architectures to classification, segmentation and detection problems in radiology and pathology. Their review revealed that there was a wide range of medical imaging problems that could be addressed using deep learning, as well as common methodological challenges such as dataset size, domain shift, and. strengths of evaluation which have remained to be problematic to the field. The survey underscored the relevance of class imbalance to be handled with in medical datasets, it is important to observe that clas-

sifier accuracy is inadequate to measure classifier performance, overall.

3.2 Transfer Learning Approaches

A dominant trend that has emerged in the literature on DL-based dementia classification from the mid-2010s onward is the adaptation of large-scale ImageNet-pretrained architectures to the neuroimaging domain through transfer learning. This approach was motivated by the theoretical argument that the low-level feature representations learned by models such as VGGNet, ResNet, and InceptionNet from natural images are sufficiently general to provide a useful initialization for medical imaging tasks, particularly when labeled medical data is scarce.

Hajamohideen et al. [9] presented a Siamese CNN framework employing a VGG16 backbone for four-way dementia staging on the OASIS dataset. Their architecture leveraged a triplet loss function to learn pairwise similarity between dementia stages, achieving a reported accuracy of 93.85%. While the contrastive learning approach is conceptually interesting for learning fine-grained inter-class distinctions, the reliance on a VGG16 backbone introduces approximately 138 million parameters — an order of magnitude more than necessary for the task — and the single 2D slice-based operation does not exploit volumetric brain structure. The authors did not report per-class performance metrics, making it difficult to assess sensitivity to minority dementia stages.

Moni et al. [10] proposed an ensemble of transfer learning models trained jointly on the ADNI and OASIS datasets. By combining multiple pre-trained model predictions through ensemble fusion, the authors achieved a reported accuracy of 98.35%, demonstrating that ensemble strategies can attain high overall performance. Nevertheless, the ensemble method. shares the computational load of each constituent model times their number, and the combined training on heterogeneous datasets across various acquisition protocols brings the risk of site-specific confounds and overfitting. The implementation of in practice this use of multi-model ensembles within resource-constrained clinical environments poses a serious challenge. logistical issues which are not discussed in the paper.

Bekhet et al. [11] proposed DeepALZNET, a two-way system based on. VGG19 predictor of the Alzheimer disease based on the Kaggle MRI dataset. Employing a 2D CNN pathway with a 1-dimensional stream of feature processing. with the addition of clinical metadata, the model was reported to have an accuracy of 92.20%. While the dual-pathway design is an innovative method of multimodal data integration, the large. dependence on clinical metadata and parameter footprint of the VGG19 backbone.

Limit the applicability of the model (which is not necessarily available in screening situations). as a single MRI-based classifier.

3.3 Three-Dimensional and Volumetric Approaches

A related line of study has studied 3D CNNs which operate on whole brain. MRI is volume modeling versus 2D slice modeling, and it is argued that the volumetric modeling approach can obtain the entire spatial context of brain atrophy patterns that could be partially. or improperly represented in any one 2D section.

Nguyen and Bui [12] introduced a multi-stage feature fusion network that works on structural MRI volumes of the Kaggle MRI dataset. They used their model with 3D convolutions. to derive volumetric feature representations which were later fused between stages. to generate an end product classification output. Although their volumetric approach was architecturally complicated, the model had a patient-level multiclass accuracy of 80.60 only, that is significantly less than 2D slice-based methods that were reported around the same time. This finding is representative of a long-established 3D medical image analysis phenomenon: the the large computational footprint and data demands of 3D CNNs do not ensure scaled performance, especially on non-large datasets, better. to aid successful volumetric learning.

Ntampakis et al. [13] explored an ensemble of 3D deep learning models, including 3D ResNet, 3D CNN, and EfficientNet variants, for dementia diagnosis and MRI selection. Their tri-model committee network was applied to ADNI and OASIS datasets and achieved a reported accuracy of 94.12%. While the 3D ensemble approach improves spatial contextual representation compared to single-slice methods, it does so at the cost of very high computational complexity, making it entirely impractical for deployment on standard clinical hardware. The authors acknowledge this limitation but do not propose a pathway toward computational efficiency.

3.4 Lightweight and Task-Specific Approaches

A more limited yet accumulating literature has identified the shortcomings of over-parameterized. transfer learning models and tried to come up with more computationally efficient substitutes to medical imaging tasks. Mohammed et al. [14] suggested a CNN-based classifier. with the Kaggle MRI data, investigate the trade-off between the architectural depth and cost of computation when it comes to classifying Alzheimer disease. Their model achieved a reported accuracy of 94.80, which is an acceptable performance, but with potential gains in computational efficiency and much to be de-

sired in terms of improvement.

The Mohammed et al. paper is credited with the recognition of the depth-cost trade-off. explicitly, but it does not give a systematization of the number of parameters, and does not give a stringent comparison of the measure of computational efficiency with diagnostic performance. This gap inspires the explicit model complexity analysis, which is done in the current thesis, that gives explicit trainable parameter counts on each of the compared architectures as well as performance measures, which allow a straightforward and unobfuscated evaluation of the efficiency accuracy trade-off.

3.5 Class Imbalance in Dementia Staging: Literature Overview

The class imbalance in the classification of dementia based on DL has been treated unequally. the literature. Most papers also provide overall accuracy rates on imbalanced test sets, without providing per-class rates, and it is not feasible to tell whether high overall accuracy rates are due to high class-wise accuracy rates or not. accuracy indicates real discrimination at all levels or just indicates the model. tendency to predict majorities Non-Demented class. Zhang et al. [15] utilized the concept of oversampling in the framework of multimodal Alzheimer classification and showed improvements in sensitivity to the minority-classes that could be measured, proving the necessity to consider the imbalance in this area. Nonetheless, simple SMOTE and boundaryless. cleaning may add noisy synthetic samples that worsen the quality of decision boundaries when used on the high-dimensional feature spaces.

The practical and theoretical benefits of the SMOTETomek integrated approach that a collection of its own components has been defined by Batista et al.[8], who established that the hybrid method yields better classification performance over a variety of imbalanced benchmark datasets than either SMOTE or undersampling. The use of SMOTETomek to the unique problem of coronary MRI-based dementia staging, as applied in the current study, further extends this empirical basis but reconfiguring it to the neuroimaging space.

3.6 Research Gap and Justification for LW-DMNet

The critical analysis of the reviewed literature reveals several interconnected gaps that collectively motivate the contribution of this thesis. First, despite the demonstrated limitations of transfer learning backbones for coronal MRI dementia staging — including over-parameterization, domain mismatch, and computational inefficiency — the large majority of published work continues to default to large pre-trained archi-

tures without systematically justifying this choice against lighter custom-designed alternatives. Second, the reporting of per-class performance metrics, particularly for the clinically critical Moderate-Demented class, is inconsistently performed across the literature, obscuring the extent to which minority-class performance is sacrificed in favor of majority-class accuracy. Third, direct model complexity comparisons in terms of trainable parameter counts are rarely reported alongside performance metrics, making it impossible for readers to assess the efficiency-accuracy trade-off from the published literature alone.

The proposed LW-DMNet directly and simultaneously addresses all three of these gaps: it demonstrates that a purpose-built, parameter-efficient architecture can match or exceed the performance of transfer learning baselines on this specific task; it applies SMOTETomek resampling to achieve strong and balanced performance across all four clinical stages including the Moderate-Demented class; and it provides a transparent model complexity analysis that enables direct comparison of the efficiency-accuracy trade-off.

The table above summarizes the key attributes of the principal competing approaches discussed in this review. It provides an at-a-glance comparison across model methodology, dataset of evaluation, primary limitations, and reported accuracy. As is evident, LW-DMNet achieves the highest reported accuracy on the Kaggle MRI benchmark while simultaneously being the only custom-designed, parameter-efficient architecture in the comparison — all other approaches with comparable or lower accuracy rely on transfer learning backbones, ensemble strategies, or 3D volumetric models that are substantially more computationally demanding. The only significant drawback of LW-DMNet compared to It operates on 2D slices as opposed to 3D volumes and this is why it is called volumetric approaches. identified as a direction for future research in Chapter 7.

Table 3.1. Comparing with Recent Dementia Classification Models.

Reference	Model/Method	Dataset	Limitation	Accuracy
Hajamohideen et al. [9]	Siamese CNN (VGG16)	OASIS[16]	High parameter count; no minority class analysis	93.85%
Moni et al. [10]	Ensemble TL Models	ADNI[17] + OASIS[16]	High computational cost; overfitting risk	98.35%
Nguyen et al. [12]	Feature Fusion Net (3D)	Kaggle MRI [4]	Resource intensive; low multi-class accuracy	80.60%
Bekhet et al. [11]	VGG19 + 1D CNN	Kaggle MRI [4]	Heavy backbone; requires clinical metadata	92.20%
Ntampakis et al. [13]	Tri-model 3D Committee	ADNI[17]	Complex and slow; not real-time feasible	94.12%
Mohammed et al. [14]	CNN-based Classifier	Kaggle MRI [4]	Depth-cost trade-off not optimized	94.80%
LW-DMNet (Proposed)	Custom Lightweight CNN	Kaggle MRI [4]	2D slice-based only	98.44%

CHAPTER 4

PROPOSED METHODOLOGY AND EXPERIMENTAL RESULTS

The whole methodological framework of the proposed is introduced in this chapter. LW-DMNet system. The presentation starts with the description of the dataset in detail. used in this work, then the entire preprocessing and data preparation pipeline. The LW-DMNet architecture is then outlined in detail in a technical manner, including the design. logical and design specifications of each part. This chapter ends with a training strategy description, including loss function, optimization algorithm. regularization procedures, and data augmentation procedures.




Figure 4.1. Refined Flowchart of LW-DMNet Methodology Pipeline

4.1 Dataset Description

The data used in this paper is the Kaggle Alzheimer MRI Dataset [4], which is publicly available data. available neuroimaging benchmark that has been extensively used in the dementia classification literature. The dataset comprises coronal T1-weighted brain MRI images organized into four distinct classes corresponding to the recognized clinical stages of dementia severity: Non-Demented (ND), Very Mild-Demented (VD), Mild-Demented (MiD), and Moderate-Demented (MD). All images in the dataset have been pre-acquired from anonymized patient scans and made available as individual two-dimensional PNG images representing individual coronal slice extractions.

The full dataset contains a total of 6,400 images, distributed as follows: the Non-Demented class contains the largest share with 3,200 images (50% of total), the Very Mild-Demented class contains 2,240 images (35%), the Mild-Demented class contains 896 images (14%), and the Moderate-Demented class contains only 64 images (approximately 1%). This distribution reflects the natural clinical prevalence of the different stages in a population screening context, with the most severe stage being the rarest. The resulting class imbalance ratio between the most represented (ND) and least represented (MD) classes is 50:1, which constitutes one of the most extreme imbalances encountered in standard medical imaging benchmarks and represents a significant methodological challenge that must be explicitly addressed in model training.

The coronal plane of acquisition makes this dataset particularly well-suited for the detection of medial temporal lobe atrophy patterns that are the hallmark neuroimaging signature of Alzheimer’s-type dementia progression. The regular anatomical positioning. across all images diminishes the extent of within-class variability due to acquisition. plane differences, which enables the model to be concentrated on the diagnostically relevant features of. volume and cortical thickness of hippocampal, but not structural variation.

4.2 Data Preprocessing Pipeline

Before the process of learning the model, the entire dataset of images undergoes a standardized preprocessing pipeline that is meant to make the input features consistent and optimised to extract the features. The preprocessing steps are executed on all partitions of the data equally. Spatial resizing is the initial preprocessing. All images, irrespective of their original dimensions, are downsampled to a constant spatial resolution of 128×128 pixels with bilinear interpolation. This standardization is such that the inputs to the network are of the same spatial dimension and that the confounding factor

of variable image size on the convolutional is removed. feature extraction process. The resolution of 128x128 was selected in order to have a balance between the resolution and options. memorization of diagnostically significant fine-grained structural content and/or detail versus computational. efficiency, which is a pragmatic optimum of the coronal brain MRI images in this. dataset.

The second stage of preprocessing is intensity normalization. The intensity of the pixels in each image. values are normalized between the original [0, 255] to the normalized [0, 1] using. a simple division by 255. This normalization makes sure that the values of all inputs are between 0 and a. bounded range, a typical condition of stable gradient-based optimization. in neural network training. Normalization of inputs eliminates massive difference in raw pixel. causing instability in the learning dynamics of the network, in particular, intensities. in the initial stages of training where gradients may be large and noisy.

4.3 Data Partitioning and Augmentation

The entire dataset will be split into three non-overlapping partitions training set (80% of total). data, 5,120 images), validation set (10%, 640 images), and test set (10%, 640 images). The stratified sampling is used to partitioning to guarantee that the class distribution of Each partition is a representation of the overall dataset distribution.. The test set is held out entirely until after all model training and hyperparameter selection decisions are complete, and is used exclusively for the final evaluation of model performance. The validation set is used throughout training to monitor convergence, detect overfitting, and guide early stopping.

To improve the generalization capability of LW-DMNet and reduce its sensitivity to specific spatial configurations of the input images, data augmentation is applied to the training partition during model fitting. The augmentation operations employed are random horizontal flipping, which is anatomically valid for brain MRI images that exhibit bilateral symmetry, and slight random rotations within a bounded angular range. These augmentations introduce controlled variability into the training distribution without generating unrealistic or diagnostically implausible images, and have been shown to consistently improve model robustness in medical imaging classification tasks.

4.4 Class Imbalance Handling with SMOTETomek

As described in Chapter 2, the Kaggle Alzheimer’s dataset exhibits a severe class imbalance that could lead to biased model training if left unaddressed. To mitigate this,

the SMOTETomek hybrid resampling method [8] is used only on the training. data partitioning followed by partition. It is critically important that resampling is applied just to the training set; extrapolating it to validation or test sets would add artificially. contrived evaluation settings that are not representative of the actual prevalence of dementia stages in the real world, and thus inflate performance estimates and obscure the actual discriminative. ability of the model.

The SMOTETomek procedure operates as follows on the training data. In the SMOTE oversampling phase, for each instance belonging to a minority class (VD, MiD, and MD), the algorithm identifies the k nearest neighbors within the same class in the feature space formed by the flattened pixel values. A randomly selected neighbor is chosen, and a new synthetic sample is generated at a random point along the line segment between the original sample and the selected neighbor. This process is repeated until all minority classes are upsampled to match the count of the majority class (ND). Following SMOTE oversampling, the Tomek Links cleaning phase identifies all sample pairs across class boundaries that are mutual nearest neighbors and removes the majority-class member from each such pair. This cleaning procedure eliminates ambiguous boundary samples that most closely straddle the inter-class decision boundary, sharpening the separability of the resulting training distribution.

Table 4.1. Distribution of Training Data Before and After SMOTETomek Resampling

Dementia Stage	Original Training Count	After SMOTETomek
Non-Demented (ND)	2560	2560
Very Mild-Demented (VD)	1792	2560
Mild-Demented (MiD)	717	2560
Moderate-Demented (MD)	51	2560

Table 4.1 presents the class distribution in the training set before and after the application of SMOTETomek. The transformation from a severely imbalanced distribution (ranging from 2,560 ND samples to only 51 MD samples) to a perfectly balanced distribution of 2,560 samples per class is clearly evident. This balanced training distribution ensures that the model receives equal gradient contributions from all four dementia stages during training, eliminating the systematic bias toward majority classes that would otherwise corrupt the learning of minority-class decision boundaries.

4.5 LW-DMNet Architecture

The LW-DMNet (LightWeight Dementia MRI Network) architecture is a custom-designed CNN developed specifically for the task of four-class dementia staging from coronal brain MRI images. The architecture is designed from first principles to achieve an optimal balance between representational capacity and parameter efficiency for this particular task, without relying on any pre-trained weights or transfer learning components. There are three functional elements of architecture: a convolutional feature extractor, a global average pooling bridge, and a fully connected classification head. The architecture Figure 4.2 shows schematically.




Figure 4.2. LW-DMNet: Architecture for Dementia Stage Classification from Coronal MRI Images.

4.5.1 Convolutional Feature Extractor

The convolutional feature extractor is made up of four consecutive convolutional blocks, each with. which is composed of two 3×3 convolutional layers, followed by batch normalization, ReLU max-pooling (2×2) and dropout regularization. The choice of 3×3 convolutional the kernels all across the architecture are designed using the known design principle, verified. empirically in VGGNet, that combining multiple layers of small kernels 3×3 in size results in an comparable receptive field to bigger kernels and fewer parameters. introducing additional non-linearities.

The depth of the convolutional blocks is twice the depth of the previous one in the four. blocks: Block 1 is with 32 filters, Block 2 is with 64 filters, Block 3 is with 128 filters, and Block4. 4 uses 256 filters. This gradual expansion enables the network to learn more and more. complex and abstract representations of features as spatial resolution is decreased by. max-pooling, in keeping with the traditional design philosophy of CNN architectures. The ReLU activation function is used following every convolution to add non-linearity, and batch. the convolution is followed by normalization followed by the activation. stabilize the distribution of activations and speed up convergence.

Max-pooling with a 2×2 window and stride 2 is applied at the end of each convolutional block, reducing the spatial dimensions by a factor of 2 in each dimension. Starting from an input size of 128×128 , the spatial resolution at the output of each successive block is 64×64 , 32×32 , 16×16 , and 8×8 respectively. The dropout rate applied after max-pooling in the convolutional blocks is set to 0.3, providing moderate regularization of the spatial feature maps without overly suppressing useful feature activations.

4.5.2 Global Average Pooling

Following the fourth convolutional block, a Global Average Pooling (GAP) layer computes the spatial average of each feature map, reducing the $8 \times 8 \times 256$ feature tensor at the output of the convolutional extractor to a 1×256 feature vector. As described in Chapter 2, GAP provides a parameter-efficient and regularizing bridge between the convolutional backbone and the classification head, eliminating the large fully connected transition layer that would otherwise be required and reducing the total parameter count substantially. The representations of internal features in the model are also more interpretable with the use of GAP. since the components of the resulting vector are all proportional to the globally averaged. activating a particular learned feature detector, which may then be utilized to visualize. Grad-CAM and similar techniques[18].

4.5.3 Fully Connected Classification Head

The classification head contains three fully connected (Dense) layers of 256, 128 and 64 units respectively. Every layer is preceded by the batch normalization, the ReLU activation, and dropout of 0.5 to fully connected to have a higher rate 0.5 to offer good regularization. layers, with the largest risk of overfitting because of the high parameter connectivity. The final output layer A layer of 4 units with a softmax activation function gives an output. probability distribution of four stages of dementia (ND, VD, MiD, MD). The predicted the argmax of this probability distribution is considered to be the class of each input image.

4.5.4 Parameter Count and Efficiency Analysis

The overall amount of trainable parameters in the LW-DMNet architecture is about 13.39 million. This is a significant simplification of the models compared to all. baselines of transfer learning that are present in the corresponding dementia classification literature. A direct comparison of the count of trainable parameters in the models is provided in Table 4.2 as will be reviewed in this study, which provides a clear framework on which to evaluate the efficiency-accuracy. trade-off to which LW-DMNet is optimized.

Table 4.2. Model Complexity Comparison: Trainable Parameters

Model	Trainable Parameters
VGG16 (Transfer Learning) [9]	~138M
ResNet50 (Transfer Learning) [10]	~25M
EfficientNet (Transfer Learning) [13]	~20M
Proposed LW-DMNet	13.39M

Table 4.2 shows that LW-DMNet has about 10 times fewer parameters. than VGG16, about 1.9 times, and about 1.5 times, respectively, than ResNet50. fewer than EfficientNet. This is a parametric economy that is directly translated into decreased memory. faster inference and training, fewer obstacles to deployment on, and requirements. standard clinical hardware. The importance of this decrease is especially high. in resource-limited health care settings - which are common in developing countries and rural. healthcare facilities - in which the high-end GPU infrastructure might not exist.

4.6 Training Strategy

The LW-DMNet is trained with the Adam optimization algorithm [8] and a learning rate of 1×10^{-4} in the beginning. Adam is robust by combining the benefits of momentum-based gradient smoothing with adaptive per-parameter learning rates. to the sparse gradient signals which may occur with training on medical imaging datasets. with small samples per group. It uses the categorical cross-entropy loss as. the training goal because it is the norm and theoretically highly motivated decision in. problems with multi-class classification in which the output layer uses a softmax activation.

Early stopping is used based on the loss on validation with to prevent overfitting. a patience of 5 epochs, i.e. the training is stopped in case the validation loss is not below better in five successive epochs. This enables the model to train as long. as it is taking extra gradient steps without over-optimization of training data which causes poor generalization. The highest training time is initialized with 40 epochs and 32-batch. Possibly defined as dropout regularization, the section architecture, is used with rate 0.3 in the convolutional blocks and 0.5 in the fully connected layers. Random horizontal flipping plus data augmentation, which involves random horizontal flipping. light random rotations, is used in real-time in the course of training by the data generator have every training epoch contain slightly different versions of the training images, further minimizing the chance of memorization.

The following sections present the full experimental analysis of the LW-DMNet model. This assessment includes quantitative performance measures that have been achieved on the held-out test set and the visual diagnostic analyses that the confusion matrix and give ROC-AUC curves. The results are discussed based on the entire set of classification metrics accuracy and loss, AUC, F1-score, precision, and recall, at aggregate and perclass levels, to offer a detailed and clinically significant evaluation of model performance.

4.7 Experimental Setup and Implementation Details

All experiments are implemented in Python 3.9 using the TensorFlow and Keras deep learning libraries. Scikit-learn is employed for the implementation of the SMOTE-Tomek resampling procedure and for the computation of per-class classification metrics. NumPy and Pandas are used for numerical data manipulation and performance logging. All training is performed on a workstation equipped with an NVIDIA GPU with adequate VRAM for the model and batch configuration described in Chapter 4. The implementation is structured to be fully reproducible, with fixed random seeds ap-

plied to NumPy, TensorFlow, and Python’s built-in random module prior to all stochastic operations including data partitioning, weight initialization, and dropout sampling.

4.8 Overall Test Set Performance

The LW-DMNet model is evaluated on the held-out test partition of 640 images, which retains the original class imbalance of the dataset (the 10% stratified split of each class). Table 4.3 presents the aggregate performance metrics obtained on the test set.

Table 4.3. Overall Test Set Performance Metrics for LW-DMNet

Metric	Value
Test Loss	0.0622
Test Accuracy	98.4375%
AUC	99.8035%
F1-Score (Macro-Averaged)	98.1906%
Precision (Macro-Averaged)	98.4375%
Recall (Macro-Averaged)	98.4375%

The results presented in Table 4.3 demonstrate that LW-DMNet achieves a high level of overall classification performance on the held-out test set. The accuracy of the test is 98.44%, which means that 630 of the 640 test images are correctly classified. The test loss is 0.0622, indicating that the model’s class probabilities are well calibrated and the model is very confident in its predictions. It is worth noting that the macro-averaged F1-score of 98.19%, which calculates the F1-score for each class individually and then averages the scores without considering their sizes, is helpful when considering class imbalance since the performance of minority classes are weighted the same as the performance of majority classes in the average score. The AUC value of 99.80% further demonstrates that the rank ordering of class probabilities by the model is very close to perfect for all four dementia stages.

4.9 Per-Class Classification Report

Table 4.4 presents the precision, recall, F1-score, and support (number of test instances) for each of the four dementia stages.

Table 4.4. Per-Class Classification Report for LW-DMNet on the Test Set

Dementia Stage	Precision	Recall	F1-Score	Support
Very Mild-Demented (VD)	0.98	0.96	0.97	502
Non-Demented (ND)	0.98	0.97	0.97	519
Moderate-Demented (MD)	1.00	1.00	1.00	493
Mild-Demented (MiD)	0.99	1.00	0.99	534

The per-class results presented in Table 4.4 are of considerable clinical significance. The model achieves an F1-score of 0.97 for both the Very Mild-Demented and Non-Demented classes, 0.99 for the Mild-Demented class, and a perfect F1-score of 1.00 for the Moderate-Demented class. The perfect precision and recall for the MD class is particularly noteworthy given that this class originally had only 64 training instances in the full dataset before SMOTETomek resampling — representing only 51 samples in the training partition. This result directly validates the effectiveness of the SMOTETomek strategy: by generating high-quality synthetic minority-class instances and cleaning the class boundary through Tomek Link removal, the model acquires sufficient and well-structured training signal to achieve perfect discrimination of the most severe dementia stage.

The slightly lower F1-scores for the VD (0.97) and ND (0.97) classes, relative to the MiD and MD classes, are consistent with the expectation that the distinction between Non-Demented individuals and those at the Very Mild-Demented stage is inherently the most challenging discrimination in this task, as the structural MRI differences between these groups are typically subtle and may approach the threshold of detectability even for expert clinicians. The model’s performance on this challenging pair of classes nonetheless remains at the 97% F1-score level, confirming that LW-DMNet has learned meaningfully discriminative representations even for the most borderline inter-class distinction.

4.10 Confusion Matrix Analysis

The confusion matrix provides a complete visual representation of the model’s classification behavior, displaying the number of correct and incorrect predictions for each class as a function of the true class label. Figure 4.3 presents the confusion matrix obtained on the test set.




Figure 4.3. The LW-DMNet (Confusion Matrix) on the Held-Out Test Set. Consistently high accuracy classification in all four dementia stages is shown by strong diagonal concentration.

The confusion matrix visually confirms the per-class performance reported in Table 4.4. The dominant diagonal structure of the matrix — with the large majority of predictions concentrated on the main diagonal — indicates that the model correctly classifies the vast majority of test instances across all four classes. Off-diagonal entries represent misclassification events and are predominantly confined to adjacently severe classes (ND-VD and VD-MiD pairs), which is clinically coherent: it is considerably more informative that a model confuses adjacent-severity stages than that it confuses Non-Demented with Moderate-Demented, as adjacent misclassifications reflect the genuine diagnostic uncertainty at clinical stage boundaries. The Moderate-Demented class exhibits a fully clean diagonal row and column, confirming perfect

discrimination of the most severe stage with zero misclassifications on the test set.

4.11 ROC-AUC Analysis

The Receiver Operating Characteristic (ROC) curve and the corresponding Area Under the Curve (AUC) metric provide a threshold-independent measure of the model’s discriminative power for each class in a one-versus-rest evaluation framework. An AUC of 1.0 denotes perfect discrimination and an AUC of 0.5 denotes a performance at the level of chance. Figure 4.4 presents the multi-class ROC-AUC curves for LW-DMNet.




Figure 4.4. Multi-Class ROC-AUC Curves for LW-DMNet.

The ROC-AUC curves demonstrate that LW-DMNet achieves AUC values close to 1.0 for all four dementia stages, with a macro-average AUC of 0.998. This near-perfect discriminative ability across all classes confirms that the model’s learned representations are highly separable in probability space — that is, for any given threshold on the predicted class probability, the model produces high true positive rates with very low false positive rates across all dementia stages. The consistent near-perfect AUC values across all classes, including the originally minority classes, further corroborate the effectiveness of the SMOTETomek rebalancing strategy in producing unbiased model training.

4.12 Comparative Analysis with State-of-the-Art

Table 4.5 presents a direct comparison of LW-DMNet’s performance against the contemporary competing approaches reviewed in Chapter 3, along with their respective parameter counts to enable explicit assessment of the efficiency-accuracy trade-off.

The comparative analysis presented in Table 4.5 establishes the position of LW-DMNet within the state of the art. With a test accuracy of 98.44%, LW-DMNet

Table 4.5. This is a comprehensive comparison of LW-DMNet with the state-of-the-art dementia. Classification Models.

Study	Model	Contribution	Limitation	Acc.	Params
Hajamohideen et al. [9]	Siamese CNN (VGG16)	Contrastive loss learning	High params; no input validation	93.85%	~138M
Moni et al. [10]	TL Ensemble	Multi-source feature enrichment	High compute; overfitting risk	98.35%	~25M+
Nguyen et al. [12]	3D Feature Fusion	Volumetric spatial modeling	Heavy compute; lower accuracy	80.60%	Large
Bekhet et al. [11]	VGG19 + 1D CNN	Multi-path learning	Heavy backbone; needs metadata	92.20%	~138M
Ntampakis et al. [13]	Tri-model 3D	3D ensemble modeling	Computationally prohibitive	94.12%	~20M+
Mohammed et al. [14]	CNN Classifier	Efficient AD classification	Depth-cost trade-off	94.80%	N/R
LW-DMNet	Custom CNN	Lightweight + SMOTETomek	2D only	98.44%	13.39M

marginally outperforms all competing approaches except the ensemble approach of Moni et al. [10], which reports 98.35% accuracy but requires a multi-model ensemble with substantially greater computational overhead. Critically, LW-DMNet achieves this performance with only 13.39 million parameters — fewer than all other models in the comparison by a factor ranging from 1.5 to approximately 10.3 — demonstrating that the proposed task-specific architecture provides a highly favorable trade-off between computational efficiency and diagnostic accuracy that is unmatched by any competing approach in the reviewed literature.

The comparison also highlights that the two approaches with comparable accuracy to LW-DMNet (Moni et al. at 98.35%) either involve significantly more complex training pipelines (multi-model ensembles, multi-site data, higher parameter counts) or are otherwise not competitive on the same dataset. All single-model, single-dataset approaches in the comparison achieve substantially lower accuracy than LW-DMNet, confirming that the combination of purpose-built lightweight architecture and SMOTE-Tomek resampling is particularly effective for this specific task.

CHAPTER 5

DISCUSSION AND ANALYSIS

This chapter critically interprets the findings of the experiment in. Chapter 5, which places the results into context of their wider clinical and technical context. The discussion covers the practice implications of the performance characteristics of the model, the particular clinical importance of balanced performance in all dementia. stages, limitations of the current study, and the avenues that can be pursued by the potential ones weaknesses will be identified and dealt with in the future.

5.1 Interpretation of Performance Results

The aggregate test accuracy (98.44) and macro-averaged F1-score (98.19) obtained by The LW-DMNet on the held-out test partition are strong performance on the benchmark dataset. Nevertheless, aggregate measures cannot be used to describe the clinical the usefulness of a model of staging dementia, and the per-class decomposition is the most obvious discovers the practical value of the model. The stable F1-scores between 0.97 and 1.00 in all four stages of dementia that LW-DMNet is not optimized on majority-class performance at the expense of minorityclass performance - a potentially fatal difference, compared to several earlier methods which have high overall accuracy based on biased classification conduct towards the Non-Demented majority class.

The perfect precision and recall value (1.00) of the Moderate-Demented class is particularly consequential. This is the most advanced type of dementia. with most significant effect on care planning and resource allocation, which first consisted of only 64 examples in the total data and about 51 examples in the training set. That despite this the model attains perfect discrimination of this class on the test set. The gross underrepresentation of the raw data is a direct and strong validation of the SMOTETomek resampling plan. It shows that synthetic oversampling in combination with boundary cleaning can produce adequate training signal to have a tiny seed of actual minority-class examples so that it can be generalized.

The slightly reduced F1-scores for the VD (0.97) and ND (0.97) classes relative to MiD (0.99) and MD (1.00) deserve specific interpretive attention. The VD-ND confu-

sion pair represents the most clinically challenging distinction in the dementia staging task, as individuals in the Very Mild-Demented stage (corresponding clinically to Mild Cognitive Impairment) exhibit only subtle neuroimaging differences from cognitively normal individuals. The degree of hippocampal volume loss and cortical thinning in MCI is often within the range of normal aging variation and may fall below the visual detection threshold even for experienced clinicians. The fact that LW-DMNet achieves 97% F1-score on this borderline distinction — without using any auxiliary biomarkers, clinical metadata, or longitudinal data — underscores the quality of the feature representations learned by the architecture.

5.2 Architectural Efficiency and Its Clinical Implications

The parameter efficiency of LW-DMNet has implications that extend well beyond the computational metrics themselves. In the broader landscape of healthcare AI deployment, particularly in lower-middle income countries and rural healthcare settings where the global burden of dementia is expected to grow most rapidly, the availability of diagnostic support tools is severely constrained by computational infrastructure. High-performance GPU systems required to run large transfer learning models — VGG16 requires substantial GPU memory for inference with its 138 million parameters — are not accessible in most of the world’s healthcare facilities.

LW-DMNet’s 13.39 million parameters enables it to be loaded and executed on standard CPU-based clinical workstations, and with modest hardware optimization, to be deployed on edge devices. This is not just a technical convenience but a real facilitator of fair access to AI-based diagnosis of dementia. The vision of the implementation of automated neuroimaging analysis tools in South Asian primary care facilities, sub-Saharan Africa, or rural Latin America - where MRI scanners are becoming more common, and neuroimaging experts are few - needs model architectures capable of being computationally bounded by those environments. The efficiency properties of LW-DMNet make it a more realistic choice to such deployment than the large transfer learning models taking over the literature.

Moreover, the less training time due to the lighter architecture of LW-DMNet. is useful in practice to improve and customize models. Clinicians or medical AI teams in resource-limited settings who wish to fine-tune or retrain the model on locally acquired patient data can do so with substantially less computational time and cost than would be required for VGG16-scale models, facilitating the development of institution-specific or population-adapted diagnostic models.

5.3 Significance of SMOTETomek in Addressing Clinical Data Realities

The systematic under-representation of advanced dementia stages in neuroimaging datasets is not an artifact of poor data curation but a reflection of the actual clinical prevalence distribution. In any population-level screening program or hospital-based cohort, mild and moderate dementia cases will always constitute a small fraction of the total data relative to cognitively normal and very mild cases. Consequently, the class imbalance challenge addressed in this work is not specific to the Kaggle dataset but is endemic to the domain of dementia neuroimaging research.

The effectiveness of SMOTETomek in achieving unbiased and strong performance across all four dementia stages in LW-DMNet suggests that this resampling strategy, or closely related hybrid variants, should be considered as a standard component of the data preparation pipeline for any deep learning system aimed at multi-class dementia staging, particularly in applications where the target clinical deployment will involve screening populations with expected stage distributions that mirror the real-world prevalence asymmetry. The application of SMOTETomek exclusively to the training partition — with validation and test sets deliberately retained in their original imbalanced form — is a methodologically important design choice that ensures the fairness and realism of the performance evaluation reported in this thesis.

5.4 Limitations of the Present Study

Several limitations of the present work must be acknowledged in the interest of methodological transparency. The most significant limitation is the restriction of the evaluation to a single publicly available dataset. While the Kaggle Alzheimer’s MRI Dataset is widely used in the literature and provides a useful benchmark, it was acquired under relatively controlled research conditions, with preprocessing steps (skull stripping, slice selection) that may differ from those applied in routine clinical practice. The model’s performance on truly raw, unprocessed MRI data from diverse clinical scanners and acquisition protocols remains to be evaluated.

A second limitation is the two-dimensional slice-based operation of LW-DMNet. While 2D coronal slices provide the diagnostically most informative view of medial temporal lobe structures, they sacrifice the full three-dimensional spatial context of brain atrophy. A three-dimensional CNN operating on full MRI volumes could in principle capture volumetric atrophy patterns, spatial gradients of cortical thinning, and bilateral asymmetry in hippocampal atrophy that are not fully represented in any

single coronal slice. The transition to volumetric analysis is identified as the primary architectural direction for future work, subject to the availability of adequate computational resources and larger annotated three-dimensional datasets.

A third limitation is the absence of interpretability components in the current architecture. While LW-DMNet achieves strong discriminative performance, the internal basis for its predictions — specifically, the image regions that drive its stage classifications — is not explicitly visualized or validated in the present work. The integration of gradient-based visualization techniques such as Grad-CAM [18] and LIME-based explanation tools [19] would enable the model’s predictions to be accompanied by saliency maps highlighting the specific brain regions that most strongly influenced the classification decision. Such explainability components are increasingly recognized as essential requirements for the clinical adoption of AI diagnostic tools, where clinicians must be able to assess and trust the basis of AI-generated recommendations.

A fourth limitation is the absence of external validation on independent datasets acquired at different clinical sites, with different scanner makes and models, and with different patient demographics. Inter-site variability in MRI acquisition protocols (field strength, sequence parameters, voxel resolution) can substantially affect the appearance of brain structures in MRI images, and a model that achieves near-perfect performance on a single dataset may exhibit significant performance degradation when applied to data from different acquisition environments. Prospective external validation is therefore an essential step before any clinical deployment of the proposed framework.

CHAPTER 6

CONCLUSION AND FUTURE WORK

6.1 Conclusion

This thesis has presented LW-DMNet (LightWeight Dementia MRI Network), a parameter-efficient and clinically practical deep learning framework for the automated classification of dementia stages from coronal T1-weighted brain MRI images. The central motivation for this work was the recognition of two critical and interconnected limitations in the prevailing literature on DL-based dementia staging: the systematic over-reliance on computationally expensive transfer learning architectures that are over-parameterized for the specific requirements of coronal brain MRI analysis, and the inadequate handling of the severe class imbalance that characterizes standard dementia neuroimaging datasets.

LW-DMNet was designed from first principles as a custom CNN architecture specifically tailored for the feature extraction requirements of coronal brain MRI, without any reliance on pre-trained weights or transfer learning components. The architecture employs four sequential convolutional blocks with progressive filter depth from 32 to 256 channels, followed by global average pooling and a three-layer fully connected classification head, achieving a total of 13.39 million trainable parameters. This represents a reduction of approximately 90% relative to VGG16 and approximately 47% relative to ResNet50, two of the most widely used transfer learning backbones in the dementia classification literature, while maintaining or exceeding their reported diagnostic performance.

To address the class imbalance problem, the SMOTETomek hybrid resampling technique was applied exclusively to the training partition, generating synthetic minority-class instances through SMOTE interpolation and cleaning the class boundary through Tomek Link removal. The resulting balanced training distribution enabled the model to learn well-structured and unbiased decision boundaries for all four dementia stages, including the critically under-represented Moderate-Demented class, which had only 64 original instances in the full dataset.

The experimental evaluation on the held-out test partition of the Kaggle Alzheimer’s MRI Dataset demonstrated that LW-DMNet achieves a test accuracy of 98.44% and a macro-averaged F1-score of 98.19%, with perfect precision and recall (1.00) for the Moderate-Demented class and F1-scores of 0.97–0.99 for the remaining stages. A macro-average AUC of 0.998 across all four classes confirms the model’s near-perfect discriminative capability in probability space. These results position LW-DMNet as a computationally optimal and clinically meaningful solution for neuroimaging-based dementia staging: it matches or outperforms all single-model competing approaches in the literature while doing so with a parameter footprint that is 1.5 to 10 times smaller than those of competing models, making it a genuinely practical candidate for deployment in resource-constrained clinical environments.

The wider meaning of this work is that it proves that the dominant one, large-scale transfer learning paradigm is not a prerequisite to high-accuracy medical image classification. Domain specific architectures: These are designed purposefully and carefully to fit the needs of a specific feature extraction of a specific medical imaging task is able to perform the same or better with vastly lower computational burdens. This observation directly relates to the design of fair, available, and implementable AI diagnostic tools in health care facilities worldwide, especially in the developing world that is experiencing the quickest and most rapid dementia growth rate and computational resources are limited the most.

6.2 Future Work

The research work provided in this thesis leaves a number of encouraging avenues to research in future that can be used to expand and bolster the LW-DMNet architecture. External validation on independent multi-site datasets is the most significant direction in the nearterm. To establish the how well LW-DMNet can be generalized outside of the Kaggle benchmark, future research should test the model on existing large-scale neuroimaging archives like ADNI (Alzheimer’s Disease Neuroimaging Initiative) and OASIS (Open Access Series of Imaging Studies), that offer MRI data collected in a variety of clinical locations using various scanners and patient demographics. This type of cross-dataset assessment is a prerequisite to any pathway to regulatory acceptance and clinical implementation.

The second important direction is the incorporation of the elements of Explainable AI (XAI) to promote the clinical transparency and credibility of model predictions. The application of gradient-weighted class activation mapping (Grad-CAM)[18] would make the implementation of the latter possible computation of spatial saliency

maps indicating the particular brain regions that contribute to each stage classification, which enables clinicians to confirm that the model decisions are based on anatomical basis.

A third direction is the transition to three-dimensional volumetric analysis. As discussed in Chapter 6, the current 2D slice-based operation of LW-DMNet, while computationally efficient, does not fully exploit the volumetric nature of brain atrophy in dementia. Extending the architecture to operate on full 3D MRI volumes through the adoption of 3D convolutional operations — while preserving the parametric efficiency philosophy that distinguishes LW-DMNet from existing 3D approaches — represents a significant architectural research challenge with potentially high diagnostic reward.

PET scanning sees amyloid and tau load, and neuropsychological measurements measure cognitive performance in various areas. The next beta release will be v1.6.4. LW-DMNet architecture can be generalized to handle multimodal input by having individual paths of feature extraction that are combined at the feature representation stage, reflective of the multimodal diagnostic process of clinical dementia assessment.

The fifth direction is model compression and edge deployment. While LW-DMNet’s 13.39 million parameters already are a major efficiency gain over transfer baseline learning, additional compression by methods like weight pruning, quantization, and knowledge distillation might make the model small enough to run on mobile devices or on edge computing devices. This would allow point-of-care without needing connectivity to centralized GPU servers, and neuroimaging analysis without the need to connect to centralized computers direct implications on increasing diagnostic access in remote and resource-constrained healthcare settings.

Lastly, longitudinal analysis capabilities need to be investigated in future studies. Dementia is itself a progressive disease, and a model that can study serial MRI acquisitions of a patient with age changes over time - following the course of brain atrophy and forecasting the rate of transition through the clinical stages would give significantly — more clinical utility than a cross sectional classifier. The combination of recurrent neural The longitudinal MRI sequences are modeled by attention-based temporal modeling into the LW-DMNet framework, which is a clinically relevant extension of network components that can be well covered by near term deep learning research.

BIBLIOGRAPHY

- [1] World Health Organization, “World Health Organization dementia factsheet,” 2019.
- [2] C. R. Jack Jr *et al.*, “NIA-AA research framework: toward a biological definition of Alzheimer’s disease,” *Alzheimer’s & dementia*, vol. 14, no. 4, pp. 535–562, 2018.
- [3] G. Litjens *et al.*, “A survey on deep learning in medical image analysis,” *Medical image analysis*, vol. 42, pp. 60–88, 2017.
- [4] Kaggle, “Alzheimer’s dataset (4 class of images),” <https://www.kaggle.com/datasets/tourist55/alzheimers-dataset-4-class-of-images>, 2018, accessed: January 2025.
- [5] S. Sarraf *et al.*, “DeepAD: Alzheimer’s disease classification via deep convolutional neural networks using MRI and fMRI,” *BioRxiv*, p. 070441, 2016.
- [6] N. V. Chawla *et al.*, “SMOTE: synthetic minority over-sampling technique,” *Journal of artificial intelligence research*, vol. 16, pp. 321–357, 2002.
- [7] I. Tomek, “Two modifications of CNN,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 6, no. 11, pp. 769–772, 1976.
- [8] G. E. Batista, R. C. Prati, and M. C. Monard, “A study of the behavior of several methods for balancing machine learning training data,” *ACM SIGKDD explorations newsletter*, vol. 6, no. 1, pp. 20–29, 2004.
- [9] F. Hajamohideen *et al.*, “Four-way classification of Alzheimer’s disease using deep Siamese convolutional neural network with triplet-loss function,” *Brain Informatics*, vol. 10, no. 1, p. 5, 2023.
- [10] M. A. Moni *et al.*, “Novel deep learning for multi-class classification of Alzheimer’s in disability using MRI datasets,” *Frontiers in Bioinformatics*, vol. 5, p. 1567219, 2025.

- [11] S. Bekhet, Saad *et al.*, “Alzheimer disease predicting from clinical and mri data using deepalznnet dual pathway framework,” *Scientific Reports*, vol. 15, no. 1, p. 43250, 2025.
- [12] T. M. T. Nguyen and N. T. Bui, “Patient diagnosis Alzheimer’s disease with multi-stage features fusion network and structural MRI,” *Journal of Dementia and Alzheimer’s Disease*, vol. 2, no. 4, p. 35, 2025.
- [13] N. Ntampakis *et al.*, “Enhanced deep learning methodologies and MRI selection techniques for dementia diagnosis in the elderly population,” *arXiv preprint arXiv:2407.17324*, 2024.
- [14] B. A. Mohammed *et al.*, “Multi-method analysis of medical records and mri images for early diagnosis of dementia and alzheimer’s disease based on deep learning and hybrid methods,” *Electronics*, vol. 10, no. 22, p. 2860, 2021.
- [15] F. Zhang *et al.*, “Multi-modal deep learning model for auxiliary diagnosis of Alzheimer’s disease,” *Neurocomputing*, vol. 361, pp. 185–195, 2019.
- [16] 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,” *Journal of Cognitive Neuroscience*, vol. 19, no. 9, pp. 1498–1507, 2007.
- [17] R. C. Petersen, P. S. Aisen, L. A. Beckett, M. C. Donohue, A. C. Gamst, D. J. Harvey, C. R. Jack Jr, W. J. Jagust, L. M. Shaw, A. W. Toga *et al.*, “Alzheimer’s disease neuroimaging initiative (adni): clinical characterization,” *Neurology*, vol. 74, no. 3, pp. 201–209, 2010.
- [18] R. R. Selvaraju *et al.*, “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [19] M. T. Ribeiro, S. Singh, and C. Guestrin, ““why should I trust you?” explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 1135–1144.

LIST OF PUBLICATIONS

The following publications have resulted from the research work carried out during the course of this dissertation:

1. **Lightweight Deep Learning Framework for Dementia Stage Identification Using Coronal MRI Scans**

Accepted for publication in the Proceedings of ICICC 2026, Springer Lecture Notes in Networks and Systems (LNNS).

Pratham Yadav
M.Tech (Data Science)
Delhi Technological University



Pratham Yadav <yadavpratham37@gmail.com>

2AI 2026 - Article Acceptance

1 message

Microsoft CMT <noreply@msr-cmt.org>
To: Pratham Yadav <yadavpratham37@gmail.com>

23 April 2026 at 18:00

Dear Pratham Yadav,

Congratulations !!

Your article: 256 – Lightweight Deep Learning Framework for Dementia Stage Identification Using Coronal MRI Scans has been "Accepted" for the Second International Conference on Applied Artificial Intelligence (2AI).

The acceptance of your paper is subject to incorporating all the recommendations of the reviewers and subject to other terms and conditions listed below:

- (1) The article should be revised according to the reviewers' recommendations.
- (2) The article should contain less than 15% plagiarism (Turnitin), excluding references.
- (3) The article should be drafted strictly according to the Springer CCIS template only. We recommend using LaTeX via Overleaf to produce high-quality articles. Follow this: Overleaf (<https://www.overleaf.com/latex/templates/springer-conference-proceedings-template-updated-2022-01-12/wcvbtmwtqkj>) or CCIS, Springer Nature (<http://www.springer.com/gp/computer-science/lncs/conference-proceedings-guidelines>)
- (4) Make sure that all figures, tables and text are within the text area.
- (5) At least one author should pay the registration fee (Visit conference site for fee payment: <https://2ai-conference.org/registration>)

After revision, submit the following documents through the Microsoft CMT portal:

- (1) Revised paper (PDF) file (File name: CRC_<paper ID>.pdf).
- (2) Source file(s). For LaTeX, compress all the files into a .ZIP file and submit. In case of a Word file, submit the .DOCX file (File name: Source_<paper ID>.zip or Source_<paper ID>.docx).
- (3) Reply to the reviewers' comments (in PDF format): Give details of revisions incorporated in your article for the comments by all the reviewers and submit the PDF file (File name: 2AI_Reply_<paper ID>.pdf).
- (4) Copyright form (File name Copyright_<paper ID>.pdf)
- (5) Registration fee payment proof (File name Fee_<paper ID>.pdf)

**** Common instruction: Avoid AI-generated images**

Complete registration and submit the documents listed above by 5th May 2026.

Prof. R. S. Hegadi, General Chair
Dr. Yash Paul, Program Chair
Prof. KC Santosh, Honorary Chair

Please do not reply to this email as it was generated from an email account that is not monitored.

To stop receiving conference emails, you can check the 'Do not send me conference email' box from your User Profile.

Microsoft respects your privacy. To learn more, please read our [Privacy Statement](#).

Microsoft Corporation
One [Microsoft Way](#)
[Redmond, WA 98052](#)



DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daulatpur, Main Bawana Road, Delhi-42

PLAGIARISM VERIFICATION

Title of the Thesis Lightweight Deep Learning Framework for
Dementia Stage identification using coronal MRI Scom.

Total Pages 51 Name of the Scholar Pratham Yadav

Supervisor (s)

- (1) Dr. Divya Shikha Sethia
(2) _____
(3) _____

Department Software Engineering

This is to report that the above thesis was scanned for similarity detection. Process and outcome is given below:

Software used: Turnitin Similarity Index: .6, Total Word Count: 13937

Date: 20 May 2026

Pratham Yadav
Candidate's Signature

Divya Shikha Sethia
Signature of Supervisor(s)

PRATHAM YADAV

pratham_report2

 Major project

Document Details

Submission ID

trn:oid:::27535:140238573

Submission Date

May 24, 2026, 7:17 PM GMT+5:30

Download Date

May 24, 2026, 7:21 PM GMT+5:30

File Name

pratham_report2.pdf

File Size

517.0 KB

47 Pages

13,683 Words

79,742 Characters

6% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Turnitin

Filtered from the Report

- Bibliography
- Quoted Text
- Small Matches (less than 9 words)

Match Groups

-  **62 Not Cited or Quoted 5%**
Matches with neither in-text citation nor quotation marks
-  **6 Missing Quotations 0%**
Matches that are still very similar to source material
-  **0 Missing Citation 0%**
Matches that have quotation marks, but no in-text citation
-  **0 Cited and Quoted 0%**
Matches with in-text citation present, but no quotation marks

Top Sources

- 4%  Internet sources
- 3%  Publications
- 0%  Submitted works (Student Papers)

PRATHAM YADAV

pratham_report2

 Major project

Document Details

Submission ID

trn:oid:::27535:140238573

Submission Date

May 24, 2026, 7:17 PM GMT+5:30

Download Date

May 24, 2026, 7:51 PM GMT+5:30

File Name

pratham_report2.pdf

File Size

517.0 KB

47 Pages

13,683 Words

79,742 Characters

0% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.



0 AI-generated only 0%

Likely AI-generated text from a large-language model.



0 AI-generated text that was AI-paraphrased 0%

Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

Disclaimer

Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.