Early Alzheimer's Disease Detection

Jyothi Sanam
Data Science
Anderson College of Business and Computing
Regis University, Denver, CO, USA
jsanam@regis.edu

Abstract

The deterioration of the brain begins to damage the brain years before the disease exhibits symptoms. Earlier diagnosis would allow patients to begin treatment at the point at which it is of benefit whereas modern techniques of diagnosing the disease are diagnosed only in late stages. In this project, two categories of brain MRI that display brain shrinkage is used to build a classifier to identify the Alzheimer and Mild Cognitive Impairment disease. The system incorporates Transformers in order to combine the information of both scan types to allow the model to learn what brain areas are the most important to diagnosis. The classifier identifies four levels, which include cognitively normal, stable MCI, progressive MCI and Alzheimer disease. This is important since progressive MCI patients deteriorate to full Alzheimer in five years - they require intervention the most. The project provides trained models, attention maps that indicate the brain regions that the transformer targeted and evaluation of failure points by the system.

I. Introduction/Background

The problem is timing. The brain alterations begin many years prior to the onset of symptoms. Cognitive tests are used to identify problems; however, by the time such problems are detected, much damage is already done. Clinical trials continue failing since we are curing patients at the late stage.

MRI reveals structural alterations. These scans are analyzed by radiologists individually and judgments made. Combination automated systems would pick up patterns that humans would overlook.

This is based on MRI in this project. This model is trained to be responsive to relevant locations of the brain in both types of scans. The system divides the patients into four categories, namely cognitively normal, stable MCI, progressive MCI, and Alzheimer's disease. There is a critical difference between stable and progressive MCI -they appear to be similar at present but very different in the end results.

I will construct two-stream networks per modality, apply transformer fusion, do adversarial network training, cross-validation, as well as comparison between fusion strategies and single-modality baselines.

II. Problem Statement

Do we have any surety that using deep learning to MRI images will early identify the presence of Alzheimer and differentiate progressive and stable MCI?

Information is missing in single-modality systems. Interpretation of radiologists does not match across experts. Majority of the models are trained on single-site data and fail when used on other scanners. The imbalance in classes is deplorable– there are many more normal controls than we most need to determine with MCI cases.

This project includes the entire pipeline. Data collection implies the downloading of MRI volumes and their matching to the diagnoses. Preparation entails scan registration to match the anatomical locations, normalization between scanners and cases with missing data. The exploration also involves visualization of patterns of diagnostic groups. Modeling includes the construction of encoders on each of the modalities, transformer fusion, adversarial training, and architecture testing. Evaluation implies calculating per-class measurements, creating attention maps, and failures analysis.

Does deep learning fusion of MRI scans improve early Alzheimer's detection compared to single-modality approaches, and which brain regions does the model gives more priority for distinguishing progressive from stable MCI?

III. Related Work

Liu et al. (2020) trained 3D CNNs on ADNI MRI scans and achieved 91% accuracy for AD versus CN classification [1]. They used only structural imaging and didn't address MCI staging.

Wen et al. (2020) published in Nature Medicine using both MRI and PET with late fusion—training separate models then combining predictions [2].

Yagis et al. (2021) used Vision Transformers for AD classification on MRI but didn't incorporate PET or use attention for fusion [3].

Pan et al. (2022) proposed adversarial domain adaptation for handling multi-site MRI data [4]. They improved generalization across scanners but worked with MRI only.

Odusami et al. (2023) reviewed deep learning for AD diagnosis and noted that most papers train on single datasets and don't test cross-site generalization [5].

IV. Data Description

Data Sources:

- **ADNI Dataset** - Longitudinal MRI and PET scans from 60+ sites with confirmed diagnoses. (https://adni.loni.usc.edu/data-samples/adni-data/)

V. Expected Outcomes

I will have some trained models of single-modality and fusion methods of associated configurations. The question is whether the fusion of transformers is meaningful as compared to other simpler methods.

The combination comparison is supposed to demonstrate the best techniques. My theory is that attention assists the model to acquire complex interactions among the types of imaging.

The focus on visualization will indicate the brain areas that the model is focused on. I would forecast emphasis on areas that are known to vary in Alzheimer but the anomaly patterns may indicate that the model learned artifacts instead.

VI. Methodology

The methodological approach I will use in my project is quantitative and observational research design that focuses on applied deep learning in the analysis of medical images. The essence of the problem is to categorize the severity of Alzheimer Disease (AD) on a structural MRI scan.

Data Sources and Preprocessing I have used two main sources of primary data: the Alzheimer Disease Neuroimaging Initiative (ADNI) dataset (300 3D NIfTI volumes of 89 patients, divided into AD, Cognitively Normal [CN], and Mild Cognitive Impairment [MCI]) and an augmented Kaggle Alzheimer dataset Data2 (more than 40,000 2D MRI patches of four levels of dementia severity).

I used nibabel and concurrent.futures to multiprocess to create an automated ETL pipeline to prepare the ADNI dataset. I was standardizing the 3D volume orientations by transposing the axes to a standard shape (240, 256, 160) and zero-padding so that the same axial footprint of 256 x 256 was achieved. I selected the 80 axial slices of each volume in the middle to obtain meaningful features and reduced the computational burden by min-max normalization to make the voxel intensities fall between 0 and 1. The Kaggle 2D dataset has also been made similar and converted to grayscale and padded/resized to 256 x 256.

In order to avoid data leakage a critical bias in medical imaging I applied a stringent patient-level stratified division (60% train/20% validation/20% test). This guaranteed that non-identical slices of the same patient will not be found in the training and test sets.

This project uses Python and the PyTorch as a primary deep learning platform. The hardware acceleration was managed using CUDA (Google Colab). I have manipulated the data using pandas with numpy, calculated evaluation metrics and stratified splitting using scikit-learn, and used seaborn/matplotlib to plot the data.

My modelling was carried out in three steps.

Firstly, I trained my own 4-block Convolutional Neural Network having Batch Normalization and Adaptive Average Pooling. Although this was 83.6% accurate on the Kaggle data, it severely overfitted much smaller ADNI data (99% training and 33.3% test accuracy) which suggests the weakness of training on small medical data in general.

Second, since the small size of ADNI was to be circumvented the known architectures (ResNet18 and EfficientNet-B0) were trained on the much larger Kaggle dataset. The winning model (98.36% test accuracy) was EfficientNet-B0. These trained weights were then used on ADNI task, which is to change the first convolutional layer to take in grayscale input and to replace the classifier head.

Third, to address inherent imbalances by classes (e.g. higher numbers of CN and MCI than AD) I replaced simple Cross-Entropy with Focal Loss, with the weight of each class dynamically determined by the occurrence rate of the training. In addition, I used a slow unfreezing strategy, and I optimized the last two blocks of EfficientNet backbone and the classifier.

In order to make my findings rigorous, I tested the models on extensive measures. In addition to the usual slice-level accuracy and loss curves, I used Confusion Matrices, Classification Reports (Precision, Recall, F1-Score), and Receiver Operating Characteristic (ROC-AUC) curves to test the models.

More importantly, I used a Patient-Level Majority Voting algorithm. As one patient gives rise to 80 slices, the assessment of consensus prediction over all slices gave a clinically relevant statistic, ultimately increasing my binary classification (AD vs. CN) accuracy at the slice level of 64.38 per cent to 70.00 per cent at the patient level.

VII. DATA DESCRIPTION

This project utilizes two distinct but complementary neuroimaging datasets, requiring rigorous data governance and management protocols. Dataset Scope and Nature The main data has been collected based on the ADNI database, involving 300 2-Year 3T MRI 300 annual scans of 89 distinct subjects. The data is raw 3D NIfTI as well as a metadata CSV of demographics (Age, Sex) and diagnostic group (18 AD, 34 CN, 37 MCI).

Secondary data (Kaggle) will contain 40,384 2-dimensional JPEG images (classified by four levels of dementia (Mild, Moderate, Non, Very Mild). This dataset contains original and augmented images and I used this data only in pre-training features extraction. I designed a preprocessing pipeline to dynamically process structural inconsistencies to ensure the quality of data. I cross-tabulated the file paths with the metadata CSV and ensured that it was 100 percent matched and did not miss a single target variable. Min-max normalization was used to mitigate outliers in voxel intensities on a slice-by-slice basis. I organized the datasets in a hierarchical system by placing the prepared .npy arrays in separate train/val/test directories. The code was versioned and model weights were regularly check-pointer and saved to the cloud storage (Google Drive) when new minimum validation losses were achieved.

The application of neuroimaging data requires ethical guidelines to be followed strictly. The ADNI dataset is highly controlled; it has de-identified patient data that has been made in line with the HIPAA policy. I ensured that no protected health information (PHI) was revealed by keeping this confidentiality through only anonymized subject IDs (such as 002S0413) supplied to me by the institution.

One of the main issues in this dataset ethical and technical was the problem of algorithmic bias because the classes were imbalanced. The number of AD samples (2,720 slices) was much lower in the ADNI training split (CN- 5, 200, and MCI- 6, 240). In its unaddressed form, the model would skew its forecasts to the majority classes.

This was countered by the use of Focal Loss weighted by inverse class frequencies, which caused the network to make incorrectly assigning the minority AD class count more heavily.

VIII. CONCLUSION

To ensure the quality of the data, I used a preprocessing pipeline capable of dynamically addressing the structural inconsistencies. I compared the paths of the files with the metadata CSV and discovered

that they were similar without missing target variables. To deal with the outliers, min-max normalization of the voxel intensities was done on a slice by slice basis.

I arranged the datasets hierarchically with preprocessed npy arrays stored in train/val/test directories and provenances of the files stored in form of generated text files (train preprocessed paths.txt, etc.). The code was version managed and periodically the model weights were check-pointed and cloned to cloud storage (Google Drive) with each new minimum validation losses made.

The neuroimaging data process cannot be done without high ethical consideration. The ADNI data are highly regulated, it has de-identified patient data which is in compliance with the HIPAA regulations. I also made sure that this confidentiality was maintained by not having any data that was revealing any of the guarded health information (PHI) and all that the institution sent is the anonymised subject IDs (say $002_{S0}413$).

The primary environmental and technical issue of this dataset was the bias in algorithms due to imbalance between classes. The training division of ADNI had a very small amount of AD samples (2,720 slices) compared to CN (5,200) and MCI (6,240). This model would be biased towards the majority classes without redress.

IX. Timeline

Week 1: Download the ADNI dataset, download kaggle dataset, set up PyTorch environment with medical imaging libraries, explore data distributions and verify paired scans exist.

Week 2: Implement preprocessing pipeline (registration, normalization), process full datasets, verify quality through visual inspection of random samples.

Week 3: Build unimodal baseline models for MRI, train with k-fold cross-validation, establish baseline performance metrics.

Week 4: Implement transformer fusion architecture with separate encoders and multi-head attention module, debug dimension mismatches, verify forward pass works.

Week 5: Train fusion models with adversarial domain adaptation, tune hyperparameters (learning rate, attention heads, fusion strategy), monitor for overfitting.

Week 6: Run full k-fold evaluation, generate attention visualizations, compute per-class metrics and confusion matrices, compare fusion variants against baselines.

Week 7: Conduct error analysis on misclassified cases, write final report documenting methodology and results.

References

[1] M. Liu, F. Li, H. Yan, K. Wang, Y. Ma, L. Shen, M. Xu, A. D. N. Initiative *et al.*, "Multi-modality cascaded convolutional neural networks for alzheimer's disease diagnosis," *Neuroinformatics*, vol. 18, pp. 295–308, 2020.

[2] J. Wen, E. Thibeau-Sutre, M. Diaz-Melo, J. Samper-González, A. Routier, S. Bottani, D. Dormont, S. Durrleman, N. Burgos, and O. Colliot, "Convolutional neural networks for classification of alzheimer's disease: Overview and reproducible evaluation," *Medical Image Analysis*, vol. 63, p. 101694, 2020.

[3] E. Yagis, S. W. Atnafu, A. García Seco de Herrera, C. Marzi, R. Scheda, M. Giannelli, C. Tessa, L. Citi, and S. Diciotti, "Effect of data leakage in brain mri classification using 2d convolutional neural networks," in *Scientific Reports*, vol. 11, no. 1. Nature Publishing Group, 2021, p. 22544.

[4] Y. Pan, M. Liu, K. Xia, and D. Shen, "Multi-view separable pyramid network for ad prediction at mci stage by 18f-fdg brain pet imaging," *IEEE Transactions on Medical Imaging*, vol. 41, no. 9, pp. 2316–2326, 2022.

[5] M. Odusami, R. Maskeliūnas, R. Damaševičius, and T. Krilavičius, "Analysis of features of alzheimer's disease: Detection of early stage from functional brain changes in magnetic resonance images using a finetuned resnet18 network," *Diagnostics*, vol. 11, no. 6, p. 1071, 2021.