Early Alzheimer's Disease Detection using Deep Learning

Presented by
Jyothi Sanam

The Problem Statement

- Over 55 million people worldwide are affected by Alzheimer's disease.
- Early interventions are most effective before significant brain damage occurs.
- Subtle structural changes in MRI scans correlate with disease progression
- Human visual detection is hard; deep learning offers an automated solution.

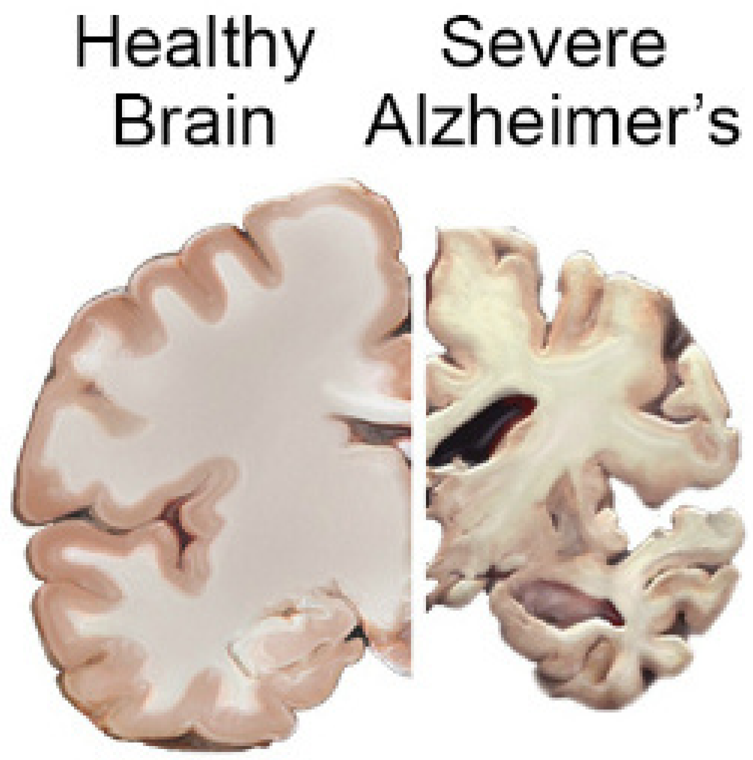

Source: National Institute on Aging (NIA), "Alzheimer's Disease Fact Sheet"

Project Objective

- Build a deep learning classifier to detect Alzheimer's from MRI images.
- The massive hurdle: Clinical datasets (like ADNI) only have dozens of patients
- Deep learning usually requires thousands of images to learn effectively.

Approach: Learn on a massive public dataset first, then "transfer" knowledge to the clinical data.

Two Tales of Data

Exploring Data Sources

- ADNI: Multi-site scanner variations and an extremely small patient count
  - https://adni.loni.usc.edu/data-samples/adni-data/
- Data2 from Kaggle: 33,984 augmented 2D slices across 4 severity classe

ADNI Alzheimer's Disease Neuroimaging Initiative

Data Insights — The Kaggle DataSet

- Severe class imbalance discovered during Exploratory Data Analysis (EDA)
- "Moderate Demented" had only 64 original images versus 3,200 "Non-Demented".
- The dataset used a 101x augmentation factor to artificially balance the rare class.
- Insight: High risk of the model memorizing artificial quirks rather than real disease features.

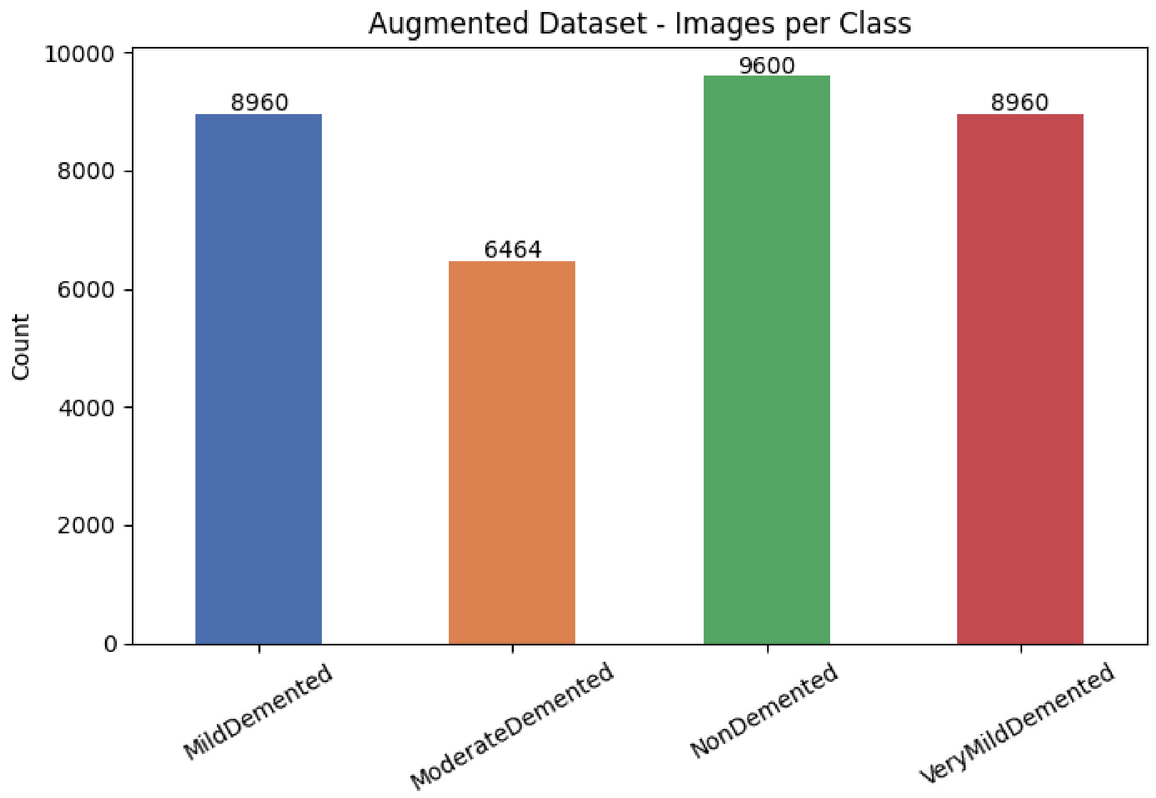

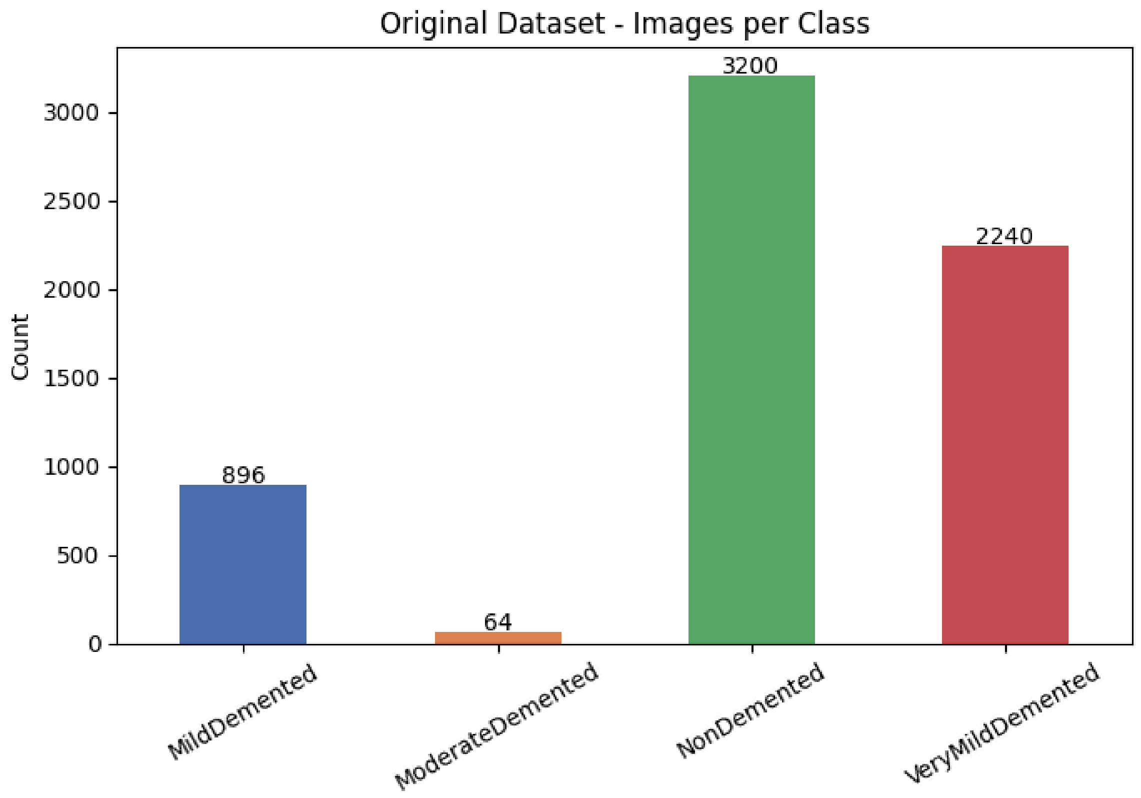

Data Insights — The ADNI Dataset

- 3D Brain scans had multiple shape variations needing alignment (e.g., transposing axes).
- MCI (Mild Cognitive Impairment) makes up 42% of patients but is clinically ambiguous.
- Age and sex were evenly distributed, ruling out basic demographic bias.

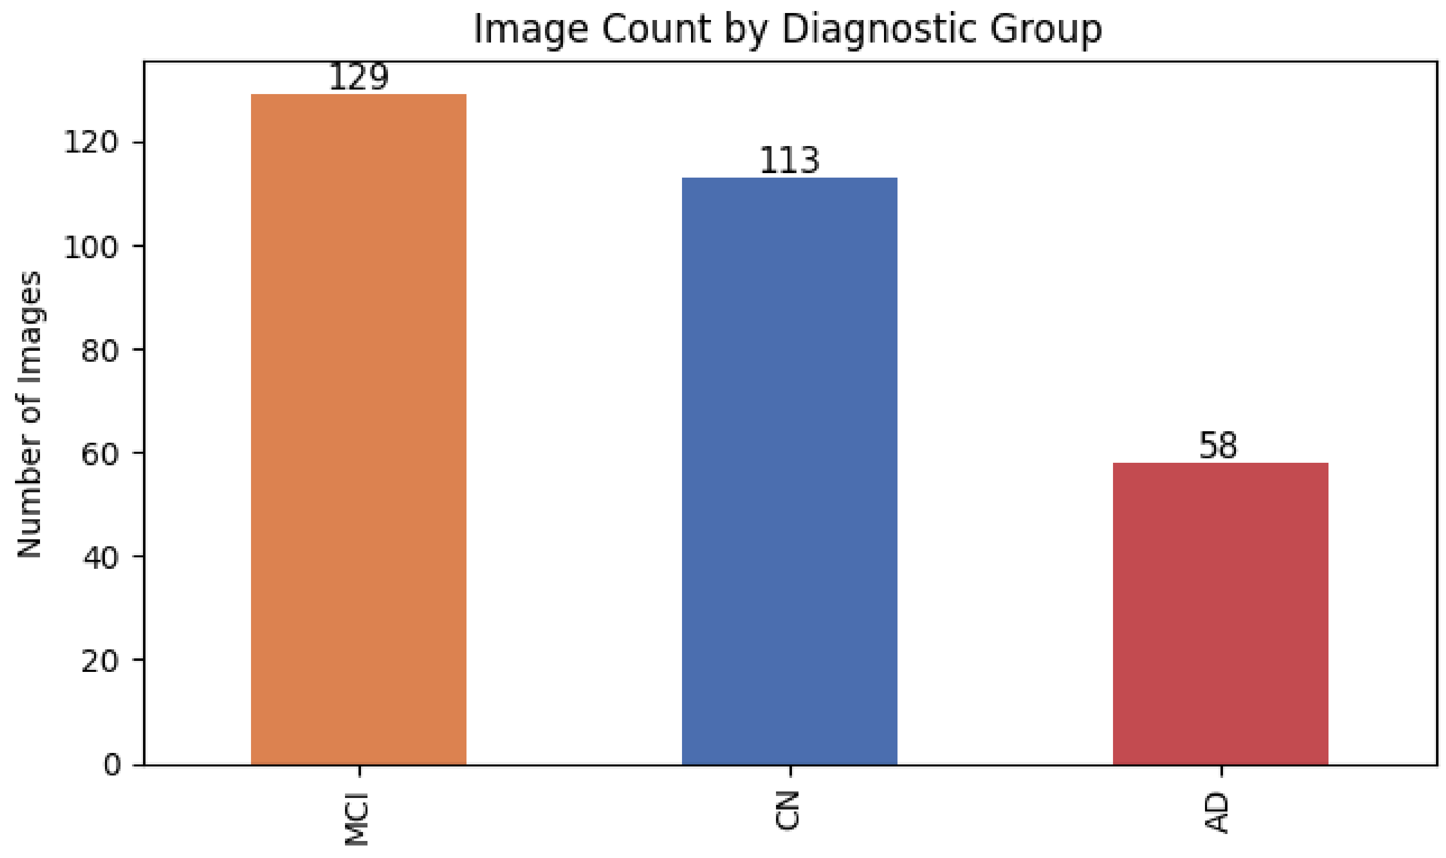

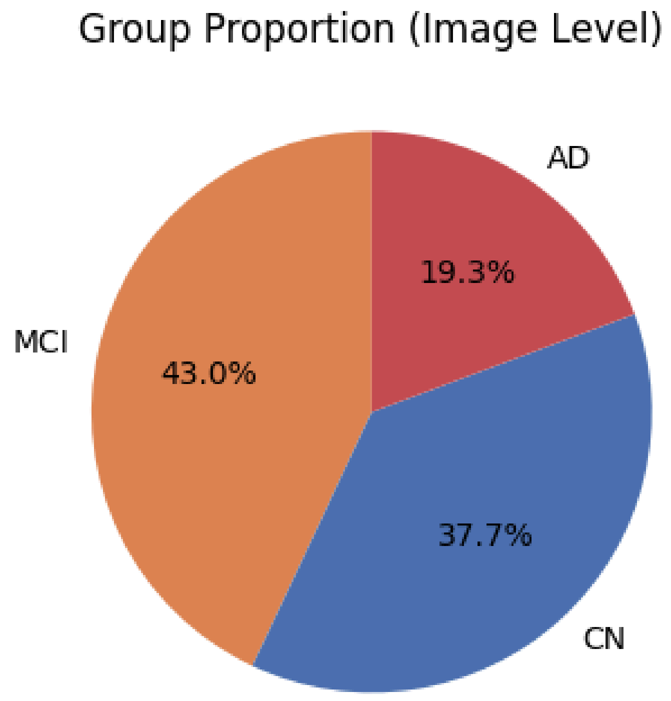

Preprocessing

- Extracted the central 80 2D slices from each 3D ADNI volume to capture key biomarkers.
- The Problem: Mixing slices randomly means slices of the *same* brain end up in both training and testing.
- The Solution: Strict Patient-Level Splitting ensures a patient's entire brain stays in one group.
- Normalization was done per-slice to account for varying MRI scanner brightness.

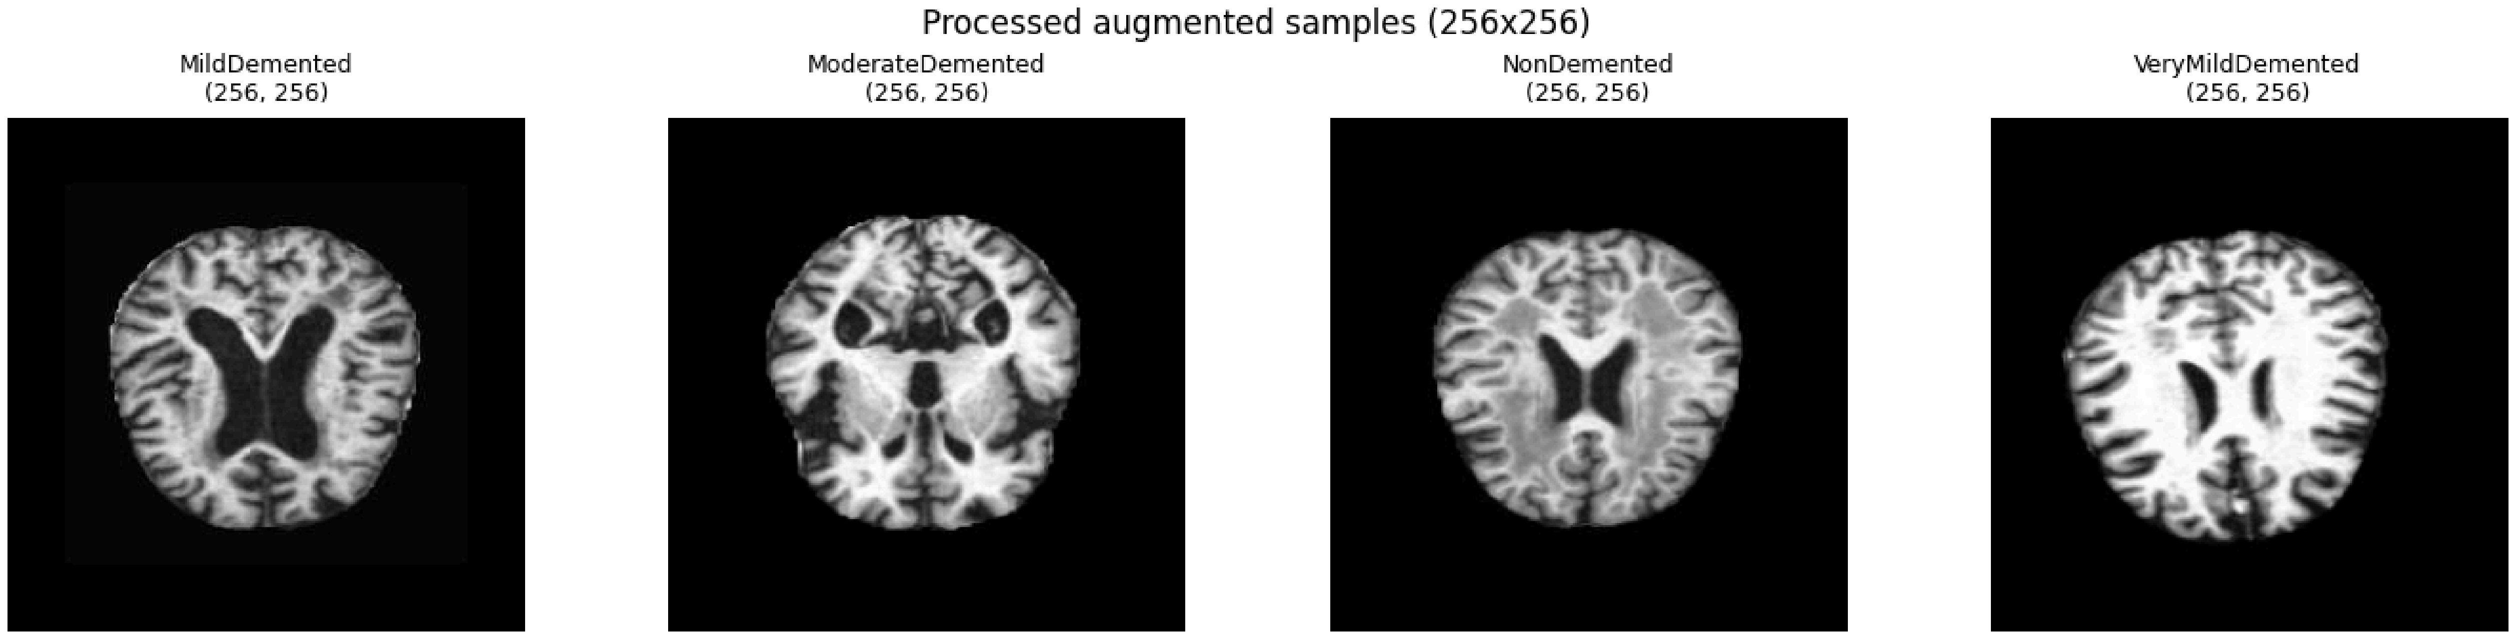

Baseline Model

- Built a custom 4-block Convolutional Neural Network (CNN).
- Used Cross-Entropy Loss and standard Adam optimizer.
- Implemented a WeightedRandomSampler to tackle the class imbalances in training batches.

Baseline Results

- On Kaggle (Data2), the baseline did well: 83.6% accuracy.
- On Clinical (ADNI), it failed entirely: 33.3% accuracy.
- The model hit its best validation score at Epoch 1 and then got worse.

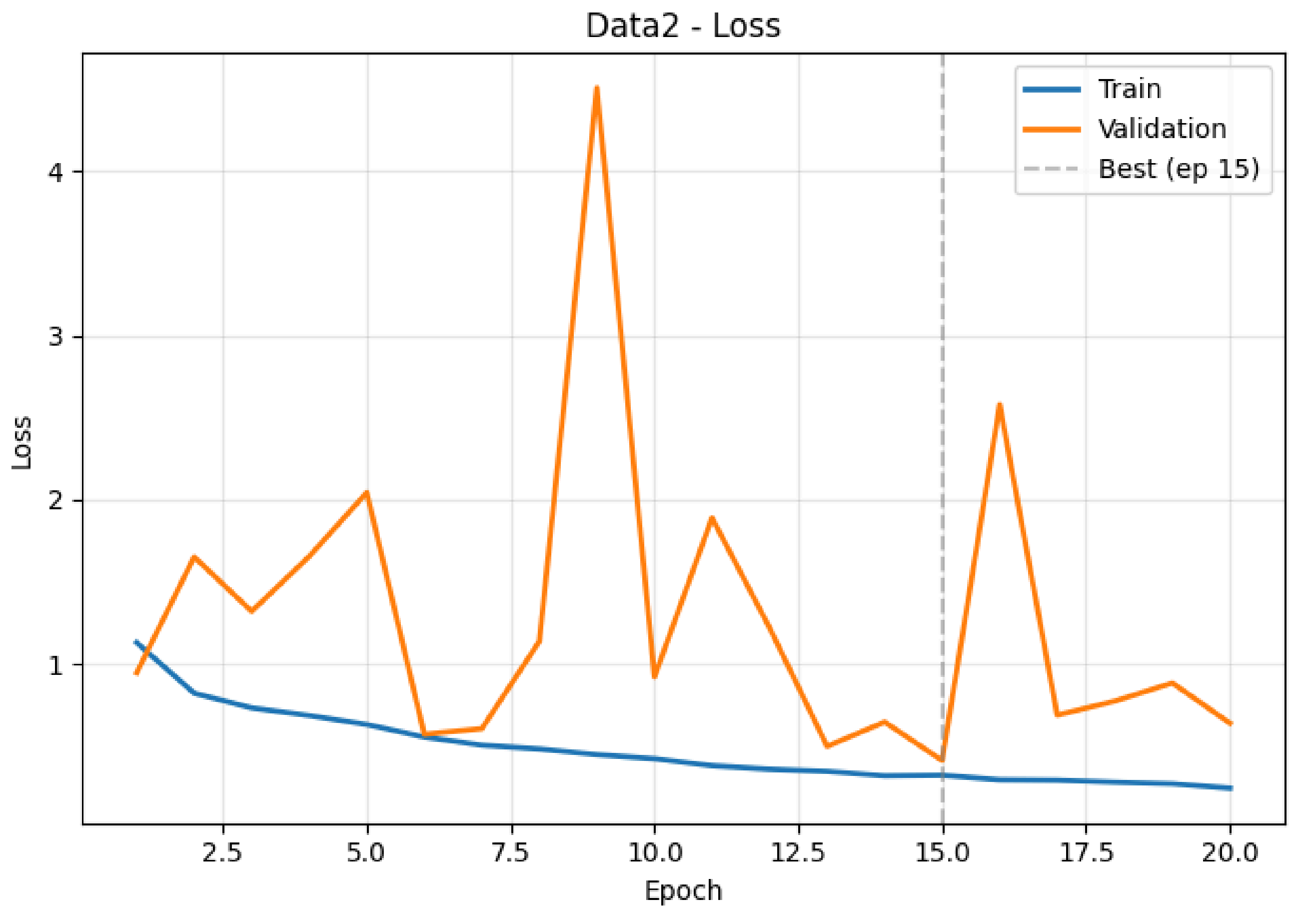

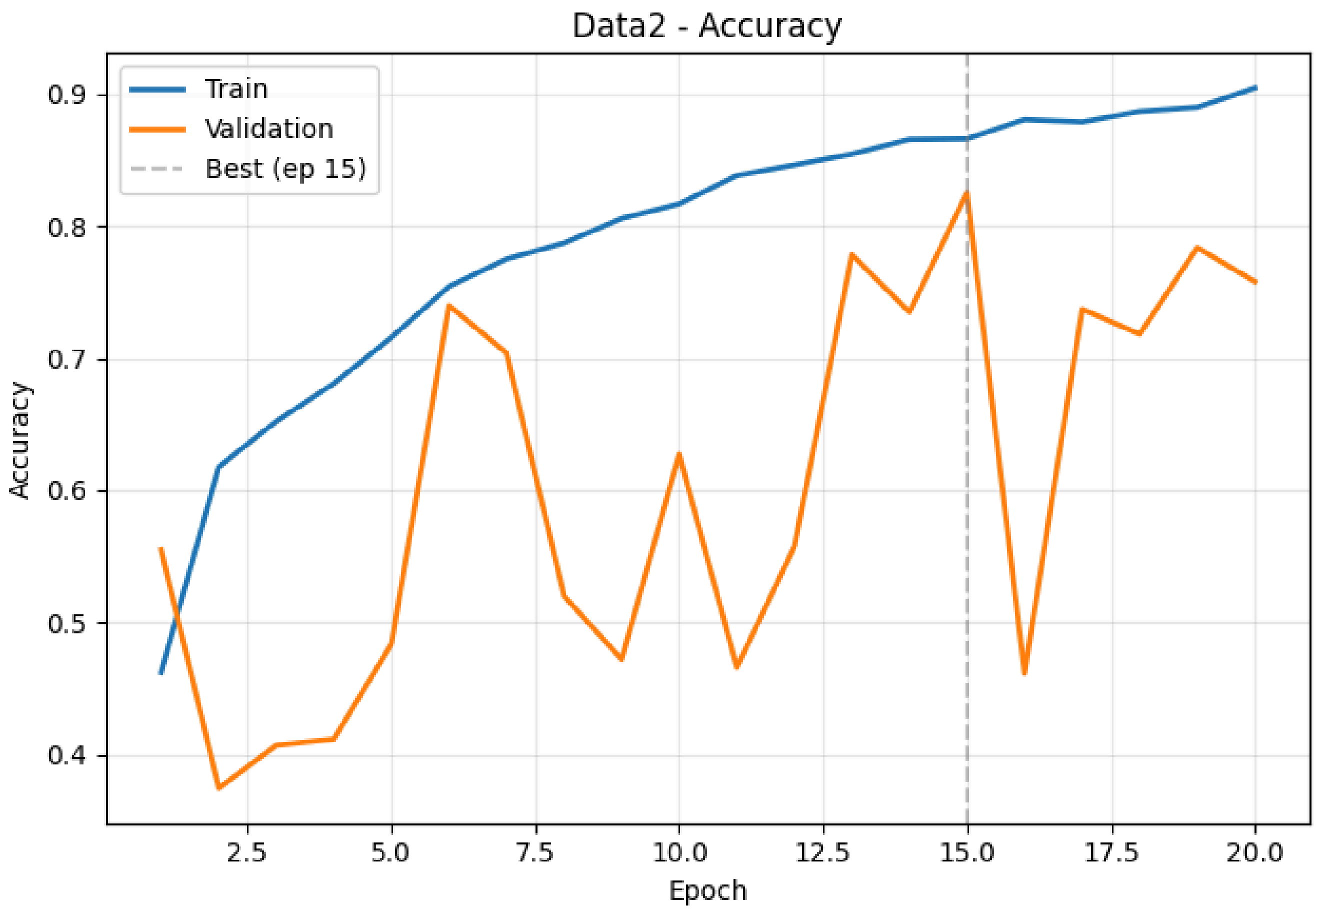

Transfer Learning

- Tested ResNet18 vs. EfficientNet-B0 on the massive Data2 Kaggle dataset.
- EfficientNet-B0 performed well.
- EfficientNet-B0 achieved a massive 98.4% accuracy on Data2.
- Near-perfect scores across all 4 severity classes.
- The model successfully learned high-quality structural patterns of brain deterioration.

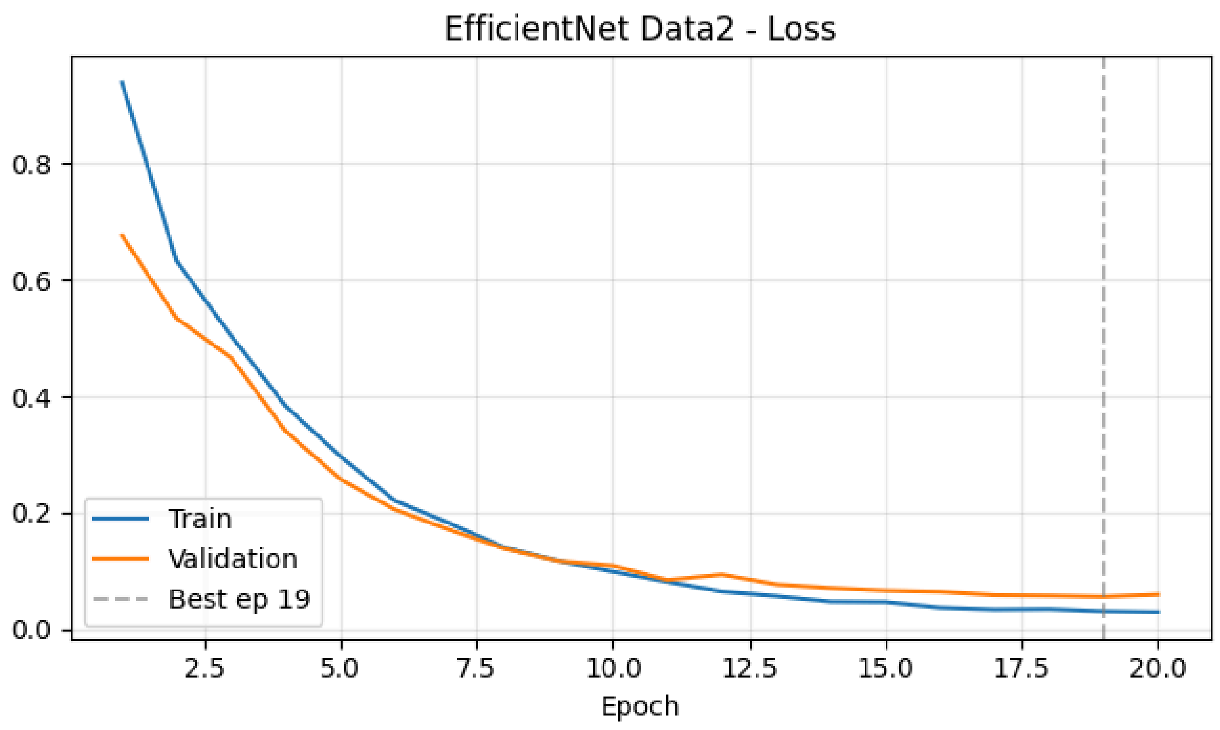

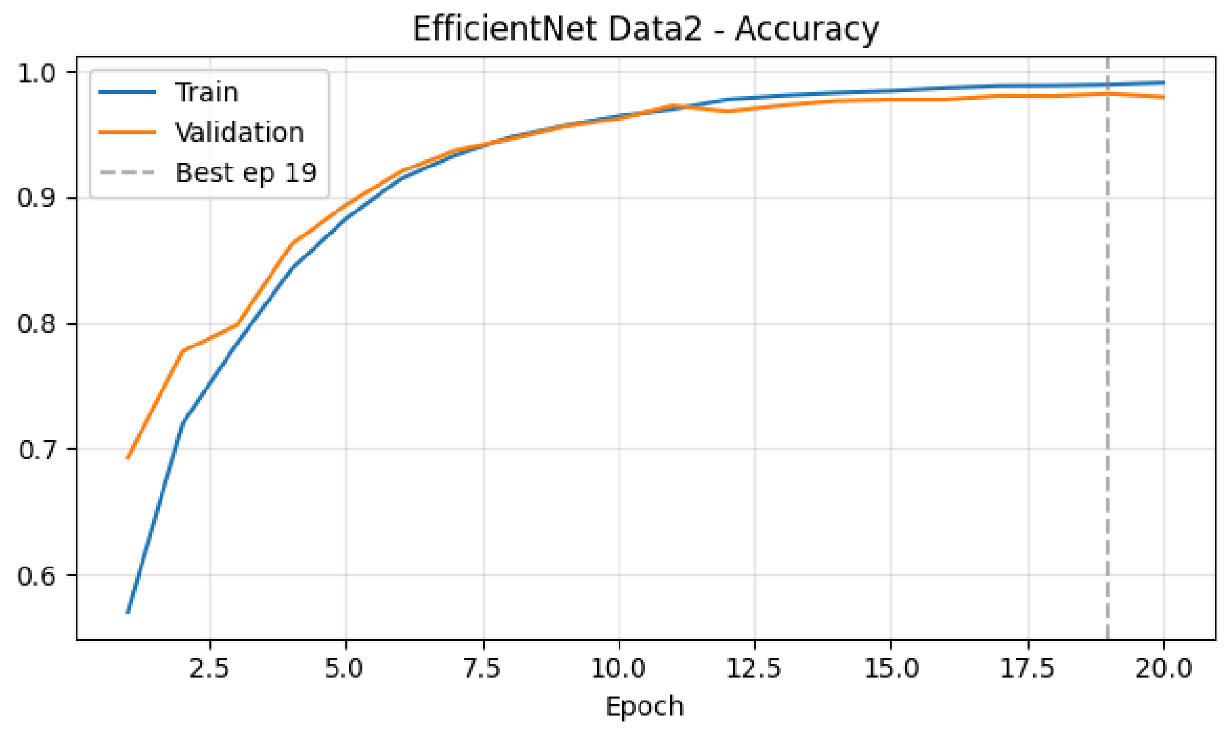

Bridging the Gap

- Loaded the 98.4% accurate Kaggle weights into the ADNI model pipeline.
- Key Decision: Frozen 90% of the model layers. Only made the final 10.4% trainable.
- Why? 89 patients are too few to train 4 million parameters without immediate overfitting.
- Shifted from a 3-class problem to a simpler, clinically standard Binary Task (AD vs. Normal)

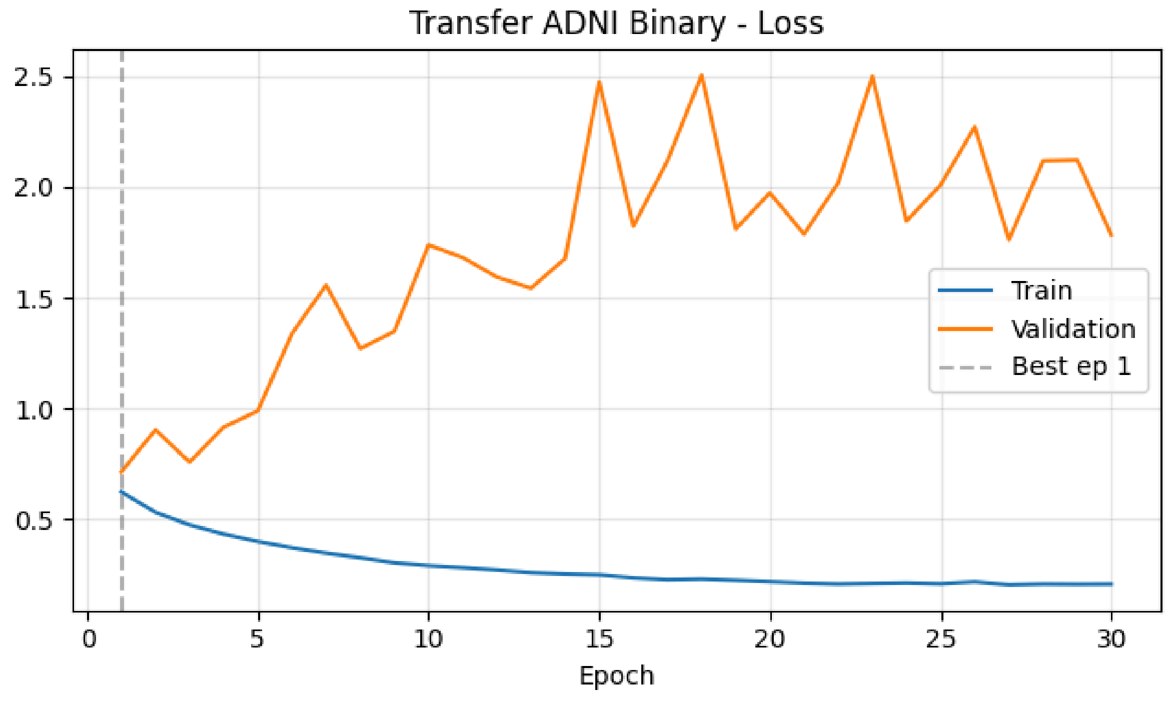

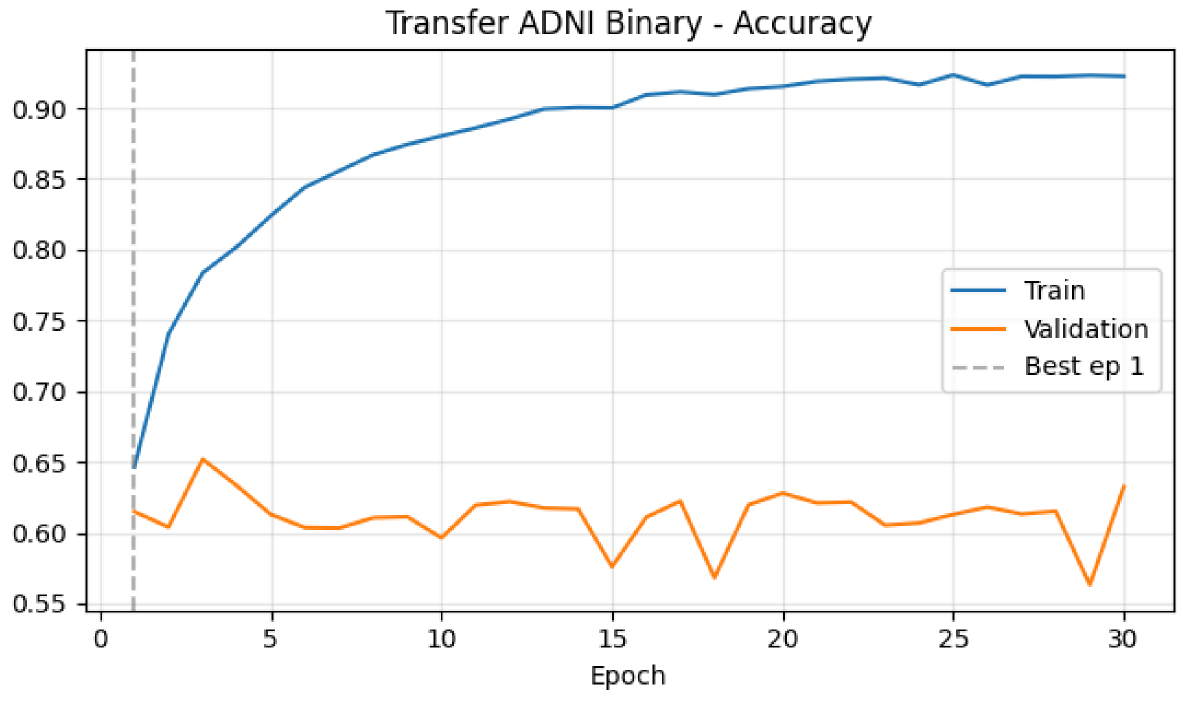

Tuning & Troubleshooting

- Despite the 70% success, the model still showed bias toward the majority class (Normal).
- Hypothesis: Standard loss functions get overwhelmed by majority classes.
- Experiment 1: Introduced Focal Loss to force the model to focus on the hardest, rarest cases.
- Experiment 2: "Unfroze" more layers (28% of the model) to adapt better to clinical scanner differences.

The Limits of Small Data

- Focal Loss *did not* improve accuracy (dropped to 30% patient accuracy).
- Deeper unfreezing recovered performance but maxed out exactly at the previous 70%.

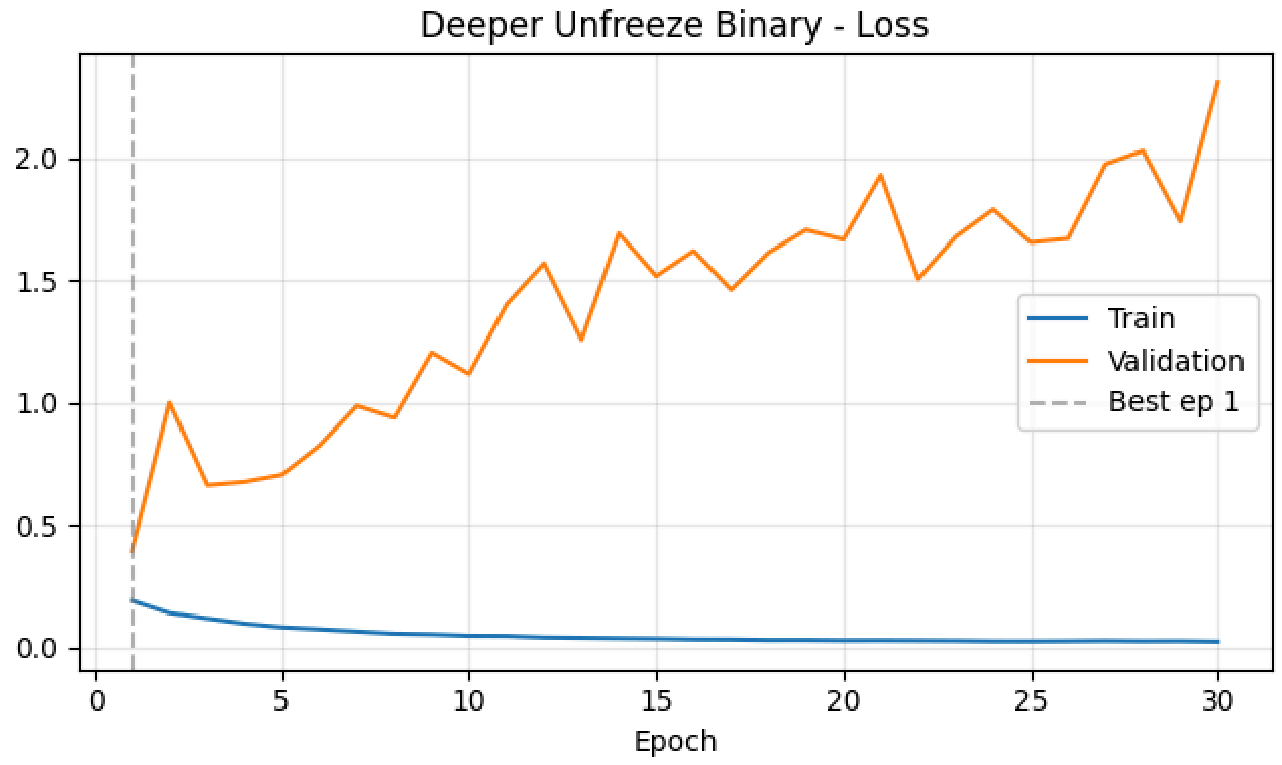

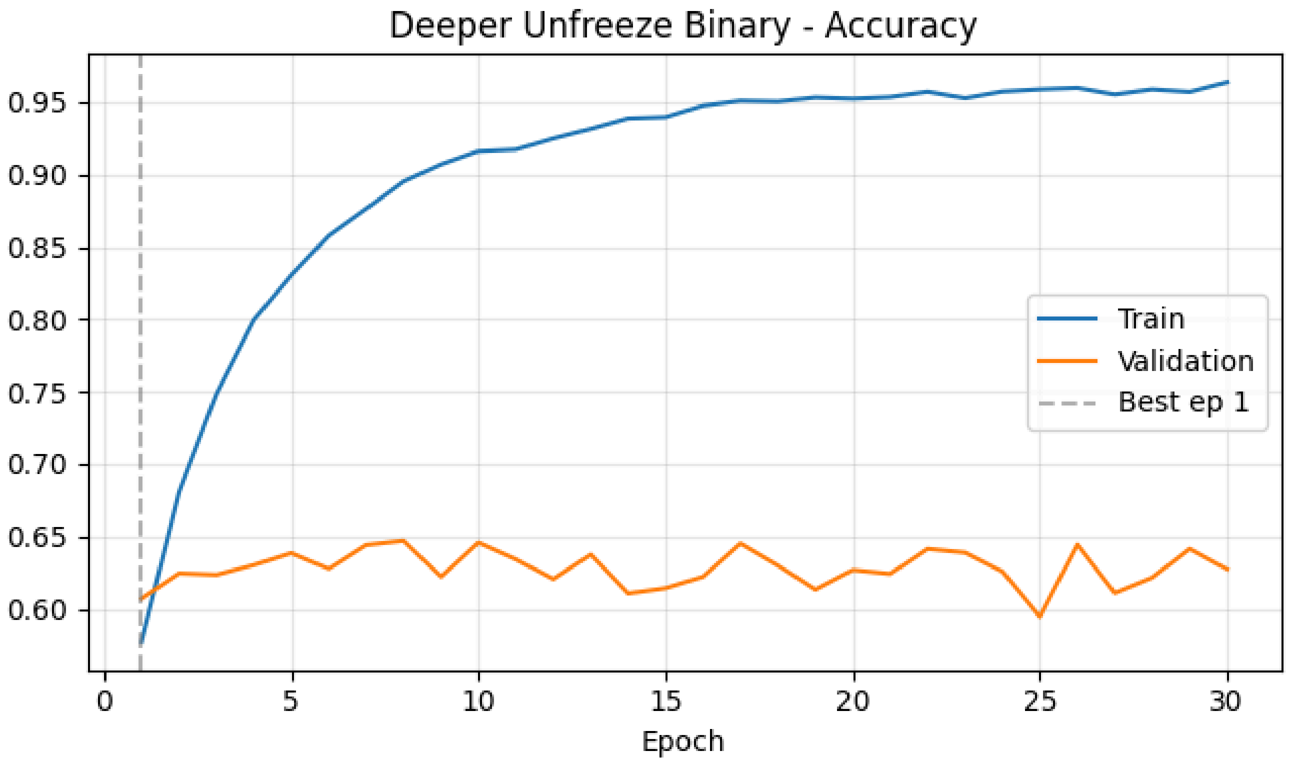

Final Results Summary

- Random Guessing Baseline: 50% (Binary).
- CNN Model: 33% (Failure).
- Final Transfer Model: 70.0% Patient-Level Accuracy.
- Published 2D slice literature achieves 65-80%; our model successfully hits the industry benchmark despite limited data.

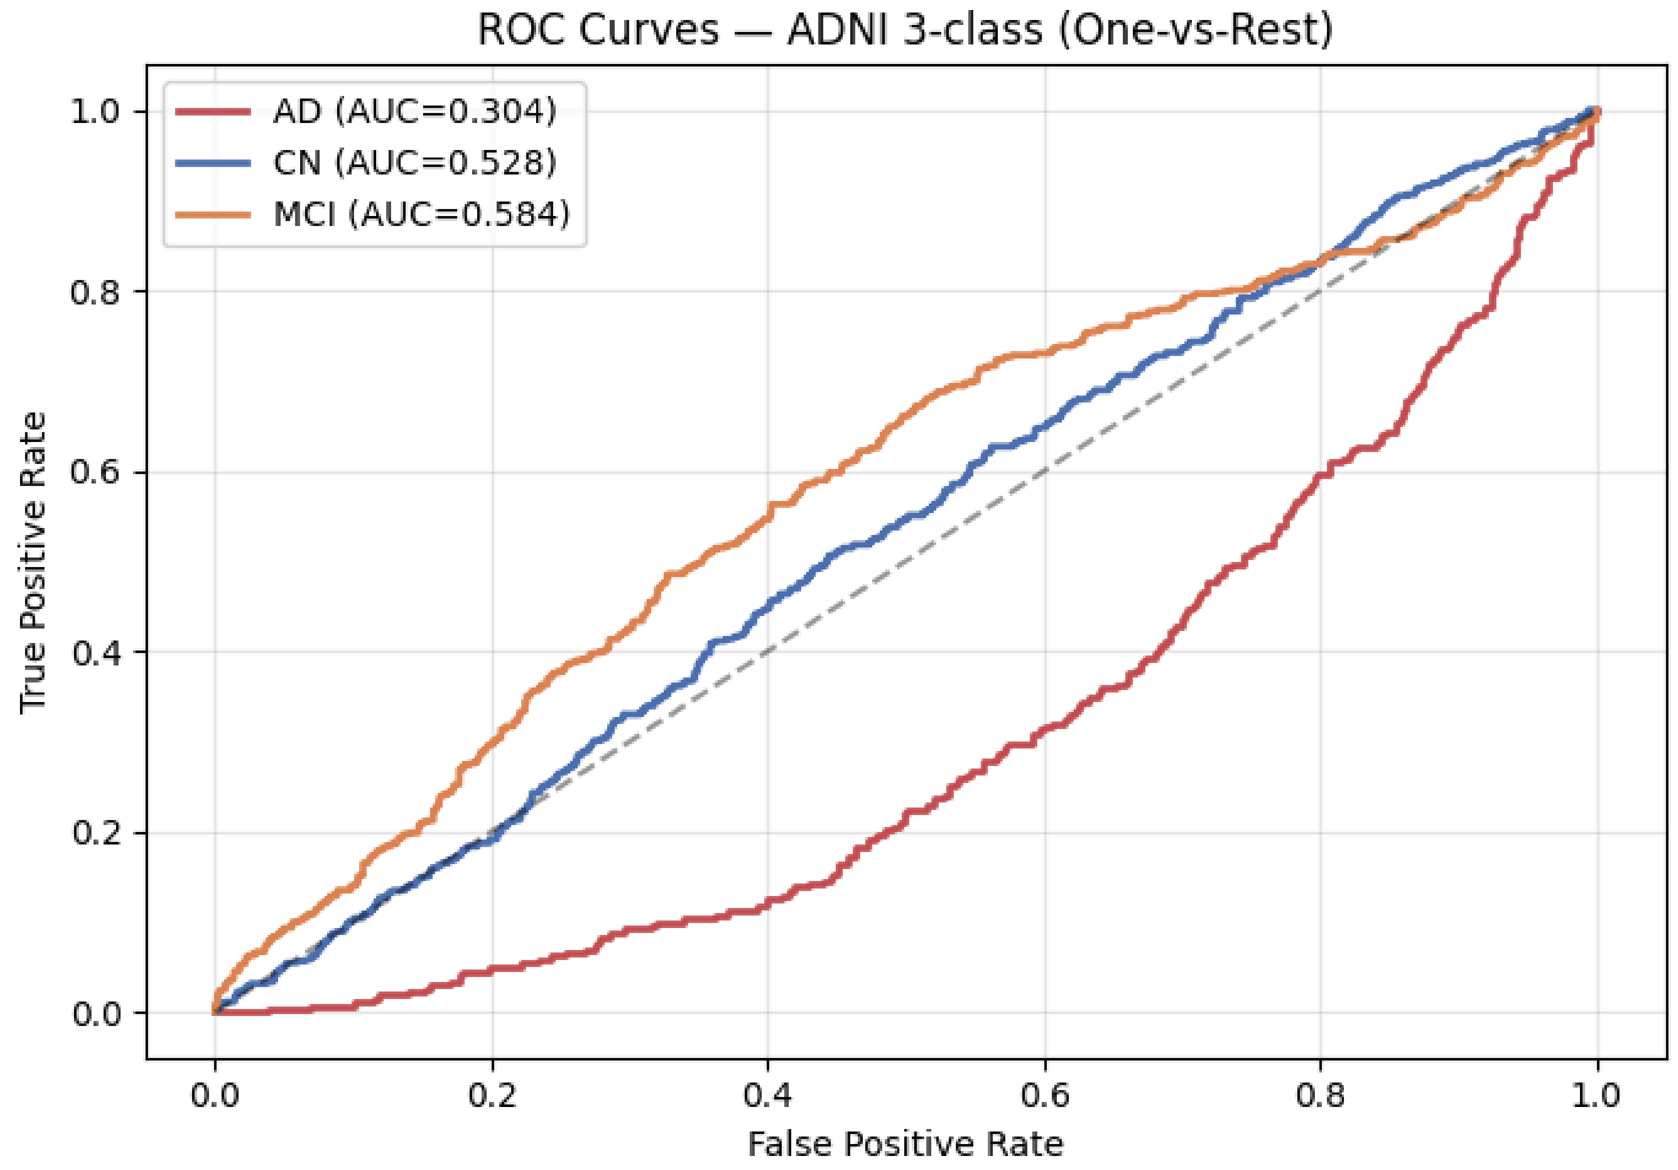

Conclusion & Future Scope

- The methodology successfully extracted meaningful clinical signals from highly constrained data.

Future Scope:

- Move from 2D slices to 3D CNNs to utilize the entire brain structure at once.
- Acquire more patient data (e.g., ADNI3) to enable robust 3-class predictions.

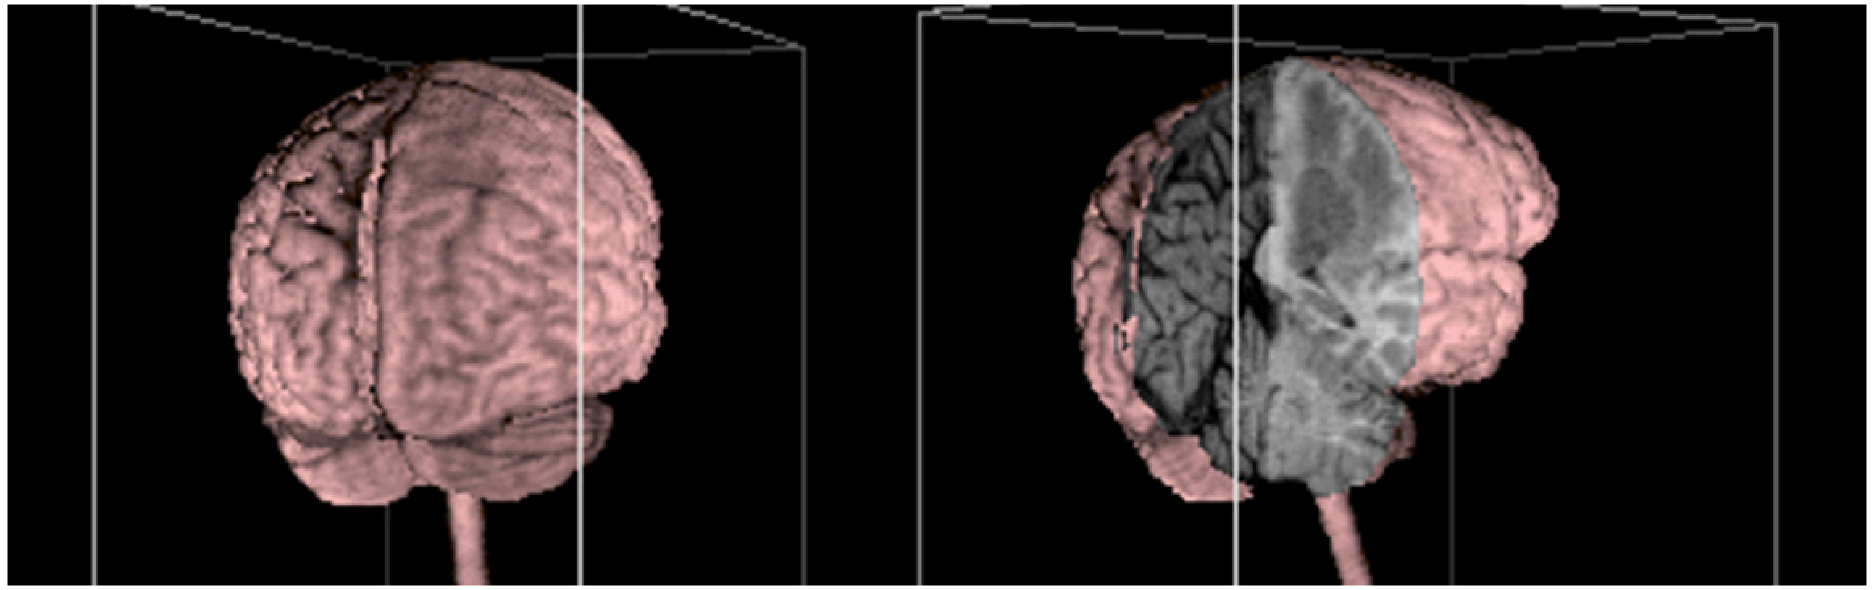

Source: The Virtual Body (VOXEL-MAN Project), 3D Reconstruction of the Brain – virtual-body.org.

Thank you