Efficient and Explainable 3D CNNs for Alzheimer's MRI Diagnosis

Martin Ropjak and Zuzana Cernekova,
Faculty of Mathematics, Physics and Informatics, Comenius University Bratislava, Slovakia

Abstract

We present a 3D ResNet-10 framework for classifying T1-weighted brain MRIs into those of cognitively normal individuals and those with Alzheimer's disease. Optimized for HPC with memory-efficient data handling and configurable parameters, it integrates Grad-CAM and occlusion for interpretability. On the ADNI dataset, the model achieved an accuracy of 91.2%, offering a scalable, reproducible, and clinically meaningful tool for medical imaging research.

Introduction & Motivation

Alzheimer's disease (AD) is a progressive neurodegenerative disorder where early diagnosis is critical. Structural MRI enables the detection of atrophy patterns, and deep CNNs achieve strong classification performance (up to 93% accuracy). Yet, many methods lack reproducibility, HPC scalability, and interpretability, the key requirements for clinical adoption.

We present an efficient and reproducible 3D CNN pipeline for AD classification, designed for HPC deployment and equipped with integrated explainability. Our main contributions are:

1. HPC-ready preprocessing and training: memory-aware chunking to optimize GPU utilization and reduce I/O overhead
2. Configurable experiments: centralized parameter control enabling easy replication and adaptation of experiments.
3. Integrated explainability with Grad-CAM variants and occlusion sensitivity to visualize salient brain regions for clinical interpretation.

We evaluate our approach on the openly available ADNI dataset at http://adni.loni.usc.edu, comparing multiple 3D CNN architectures and analyzing both classification accuracy and visual explanations

Methodology

A. Dataset

T1-weighted MRI scans from the ADNI study (openly available http://adni.loni.usc.edu).

B. Preprocessing pipeline

- Skull stripping and N4 bias field correction
- Affine registration to the MNI ICBM 152 symmetric template (1 mm isotropic, 193×229×193 voxels)
- Foreground extraction (removal of non-brain tissue)
- Intensity normalization (clipping at 99.5th percentile, rescaling to [0,1])
- Resampling to 128^3 voxels
- Converted to float16 tensors to reduce memory footprint.

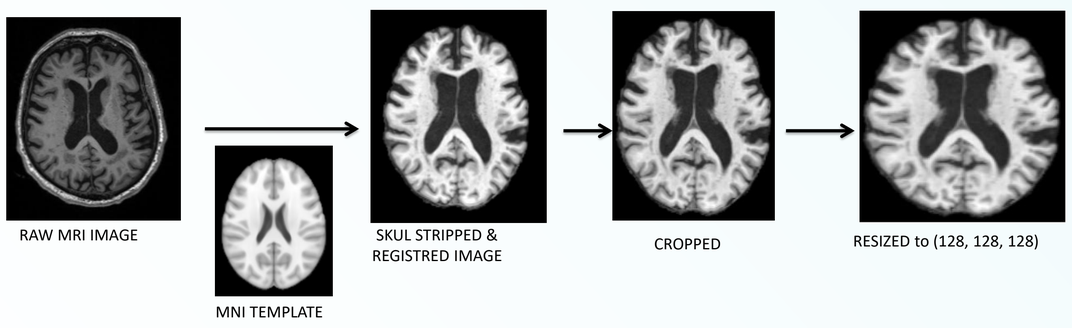

C. Memory-Aware Data Loading

- Dataset split into N preprocessed .pt chunks
- Hybrid shuffling across/within chunks for robustness
- Custom loader enables training on both high- and low-memory systems
- Stored as 5D float16 tensors → fewer disk reads & faster training

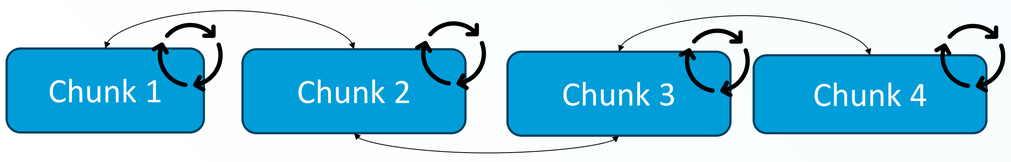

D. Model Architectures & Training

- Tested 3D ResNet, DenseNet, EfficientNet (with MedicalNet pretraining)
- Configurable dropout and unified config file for reproducibility
- Trained on Devana HPC (A100, 40GB) with Adam/AdamW, CE/Focal loss
- Batch size 8, early stopping, 3 runs per setup
- Code: github.com/MartinR255/alzheimer-disease-detection

Workflow Overview

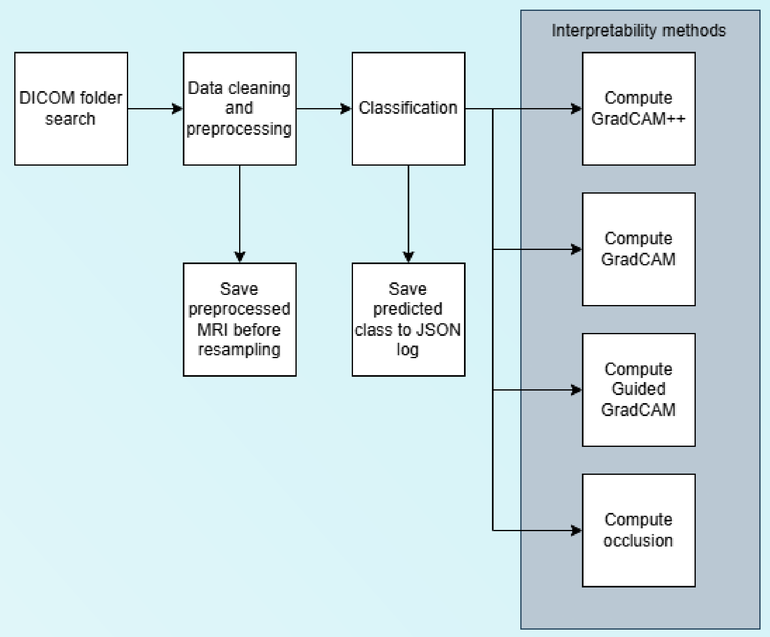

Results

Table 1: *Top 3 models by validation and test performance.*

Model	Validation Acc. (%)	Validation F1 (%)	Test Acc. (%)	Test F1 (%)
ResNet10	87.82	85.86	**91.20**	**90.09**
DenseNet121	**88.30**	86.12	90.65	89.29
ResNet18	88.06	**86.23**	88.46	87.15

Table 2: *Per-class performance of ResNet10 on the test set.*

Class	Acc. (%)	Prec. (%)	Recall (%)	F1 (%)
CN	**93.80**	93.03	93.80	93.41
AD	86.06	87.50	86.06	86.77

Explainability

Our research also addresses the critical issue of transparency and interpretability in AI models. Post-hoc methods highlight key regions in an Alzheimer's case. Grad-CAM shows structural changes consistent with AD pathology, while occlusion sensitivity confirms these regions as critical, with red areas indicating where occlusion strongly reduced model confidence.

Grad-CAM.

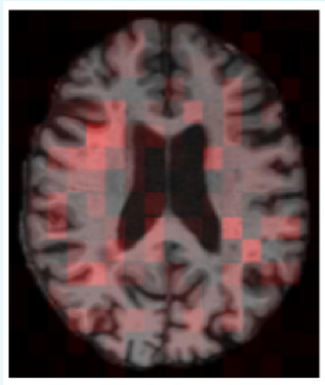

Occlusion sensitivity map.

Conclusion

Our lightweight ResNet10 achieved 91.2% test accuracy with balanced F1-scores on Alzheimer's MRI data. Memory-efficient training and a fully configurable setup ensured reproducibility, while Grad-CAM and occlusion sensitivity highlighted clinically relevant regions such as the hippocampus and temporal lobes. This efficient, explainable framework provides a solid baseline for future research.

Acknowledgements

This work was funded by the EU Next Generation EU through the Recovery and Resilience Plan for Slovakia under the project No. 09I02-03-V01-00029 and by the APVV-23-0250 project.