Eurographics Workshop on Visual Computing for Biology and Medicine (2025)
T. Gerrits and C. Gillmann (Poster Chairs)

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 configurable deep learning framework for classifying T1-weighted brain MRI scans into cognitively normal (CN), and Alzheimer's disease (AD). The pipeline is optimized for high-performance computing (HPC) environments and integrates explainability methods for visual model interpretation. Key contributions include a flexible, parameter-driven setup that enhances reproducibility, a memory-aware dataset chunking strategy for efficient GPU utilization, and built-in Grad-CAM visualizations for visual inspection of results. Using the ADNI dataset, our best binary classification (CN vs. AD) achieved 91.2% accuracy with a 3D ResNet10 model. The framework's scalability, reproducibility, and explainability make it a robust tool for medical imaging research.

CCS Concepts

• ***Computing methodologies*** → ***Computer vision;***

1. Introduction

Alzheimer's disease (AD) is a progressive neurodegenerative disorder that leads to memory loss, cognitive decline, and functional impairment. Early diagnosis is critical for timely intervention and treatment planning. Structural magnetic resonance imaging (MRI) has emerged as a key biomarker in AD research, enabling the detection of brain atrophy patterns associated with disease progression [DB23]. Deep convolutional neural networks (CNNs) achieve strong performance in structural MRI classification tasks for AD, with reported performance reaching up to 93% accuracy in binary classification [ABL*23], [RAS*25]. However, many approaches lack reproducibility, are not optimized for large-scale high-performance computing (HPC) environments, and provide limited interpretability, factors critical 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 [JBF*08], comparing multiple 3D CNN architectures and analyzing both classification accuracy and visual explanations.

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.
This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

2. Method

2.1. Data Preprocessing

We used T1-weighted MRI scans from the ADNI dataset. Preprocessing included skull stripping, N4 bias field correction, and affine registration to the MNI ICBM 152 symmetric template (1 mm isotropic, 193×229×193 voxels) [GJB*06]. Foreground extraction removed non-brain background, followed by intensity normalization (99.5th percentile clipping, scaling to [0,1]) and resampling to 128^3 voxels. Data were converted to `float16` tensors for storage.

2.2. Memory-Aware Data Loading

To reduce training overhead, MRI volumes were preprocessed once, split into memory-sized chunks, and stored as `.pt` tensors. For large datasets, a custom loader partitions the data into N chunks based on memory limits [AMB19], applying hybrid shuffling both across and within chunks to balance randomness and disk efficiency. This enabled training on both high-memory (64 GB RAM) and constrained systems.

2.3. Model Architectures

We evaluated 3D ResNet (10, 18, 34, 50, 101 layers), DenseNet and EfficientNet variants. ResNet models could optionally use pretrained MedicalNet weights. Configurable dropout was supported either after ReLU activations or before the final fully connected layer. All training settings (e.g., architecture, optimizer, learning rate, dropout rate, spatial resolution) were defined in a single experiment configuration file, facilitating reproducibility.

2.4. Training Setup

Training was conducted on the Devana supercomputer using a single NVIDIA A100 (40 GB). Models were trained with Adam or AdamW optimizers, Cross-Entropy or Focal loss, batch size 8, and early stopping based on validation accuracy. Each experiment was repeated three times to assess stability. Our implementation can be found on https://github.com/MartinR255/alzheimer-disease-detection.

2.5. Explainability

To enhance model transparency and support clinical interpretability, we integrated four post-hoc interpretability techniques: Grad-CAM, Grad-CAM++, Guided Grad-CAM, and Occlusion Sensitivity. These methods serve complementary purposes in highlighting salient brain regions that drive model decisions.

Grad-CAM [SCD*17] and Grad-CAM++ generate class-discriminative heatmaps by backpropagating gradients to the final convolutional layers, revealing spatial patterns associated with the predicted class. Guided Grad-CAM further enhances resolution by combining gradient information with activation maps. Occlusion sensitivity offers a perturbation-based view by systematically masking input patches and measuring drops in prediction confidence, pinpointing critical regions. Visualizations were produced for both correctly and incorrectly classified cases to analyze model attention and identify possible failure points.

3. Results

Models were evaluated on a held-out ADNI test set, with results averaged over three runs. The best overall performance was achieved by a 3D ResNet-10, reaching 91.20% accuracy and a 90.09% F1-score for Alzheimer's vs. control classification. DenseNet121 and ResNet18 showed comparable validation performance (accuracy ≈88%, F1-score ≈86%) but slightly lower test scores. While DenseNet121 approached ResNet10 in accuracy (90.65%), it required longer training times, and ResNet18 performed marginally worse in both metrics. These results indicate that lighter architectures, such as ResNet10, can outperform deeper models in both efficiency and accuracy.

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

	Validation		Test	
Model	Acc. (%)	F1 (%)	Acc. (%)	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

The best-performing ResNet10 model used dropout regularization (0.3 after ReLU, 0.2 before the fully connected layer), label smoothing ($\epsilon = 0.1$), and class weights to address imbalance. Training employed the AdamW optimizer ($\eta = 1e{-}4$, weight decay $1e{-}4$) for 30 epochs, yielding the optimal model at epoch 23.

A breakdown by class (Table 2) shows the network is particularly confident in identifying cognitively normal (CN) subjects, with slightly lower recall for Alzheimer's disease (AD) cases. This imbalance aligns with the skew in the training data, partially mitigated by weighted loss.

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

To enhance interpretability, we applied post-hoc methods. Figure 1 shows the visualizations for an Alzheimer's disease case. The Grad-CAM heatmap (left) reveals that predictions are primarily driven by structural brain changes consistent with known AD pathology, such as cortical atrophy and ventricular enlargement. Some background activation remains, but it is relatively minor. The occlusion sensitivity map (right) further confirms these regions as critical, with red areas indicating where occlusion significantly reduced model confidence.

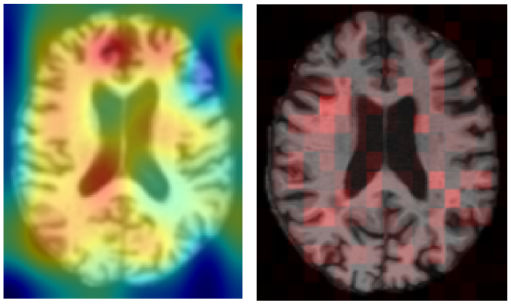

Figure 1: *Explainability visualizations for a correctly classified Alzheimer's case.* ***Left:*** *Grad-CAM.* ***Right:*** *Occlusion sensitivity map.*

4. Conclusion

Despite using a lightweight ResNet10 architecture, our approach achieved a competitive test accuracy of 91.20% and balanced F1-scores. The compact model and memory-efficient data loading enabled fast training on large 3D MRI datasets, while integrated explainability supported clinical interpretability.

The framework offers three key advantages: (1) chunked data loading reduced training time by 35%; (2) experiments are fully reproducible via a single configuration file; and (3) Grad-CAM and occlusion sensitivity consistently highlighted clinically relevant regions, notably the hippocampus and temporal lobes.

These results show that the pipeline is efficient, reproducible, and clinically meaningful, making it a solid baseline for future Alzheimer's MRI research.

Acknowledgment

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

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.

References

[ABL*23] AGARWAL D., BERBÍS M. Á., LUNA A., LIPARI V., BALLESTER J. B., DE LA TORRE-DÍEZ I.: Automated medical diagnosis of alzheimer´s disease using an efficient net convolutional neural network. *J. Med. Syst. 47*, 1 (May 2023), 57. 1

[AMB19] AIZMAN A., MALTBY G., BREUEL T.: High performance i/o for large scale deep learning. In *2019 IEEE International Conference on Big Data (Big Data)* (2019), pp. 5965–5967. doi:10.1109/BigData47090.2019.9005703. 1

[DB23] DAS S., BHARATI K.: Recent advances in treatment of alzheimer's disease. *International Journal of Alzheimer's Disease Research 1*, 2 (August 2023), 34–38. URL: https://www.opastpublishers.com/open-access-articles/recent-advances-in-treatment-of-alzheimers-disease.pdf, doi:10.9734/ijadr/2023/v1i2a34. 1

[GJB*06] GRABNER G., JANKE A. L., BUDGE M. M., SMITH D., PRUESSNER J., COLLINS D. L.: Symmetric atlasing and model based segmentation: An application to the hippocampus in older adults. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006* (Berlin, Heidelberg, 2006), Larsen R., Nielsen M., Sporring J., (Eds.), Springer Berlin Heidelberg, pp. 58–66. 1

[JBF*08] JACK C. R., BERNSTEIN M. A., FOX N. C., THOMPSON P., ALEXANDER G., HARVEY D., BOROWSKI B., BRITSON P. J., WHITWELL J. L., WARD C., ET AL.: The Alzheimer's Disease Neuroimaging Initiative (ADNI): MRI methods. *Journal of Magnetic Resonance Imaging 27*, 4 (2008), 685–691. doi:10.1002/jmri.21049. 1

[RAS*25] RAHMAN A. U., ALI S., SAQIA B., HALIM Z., AL-KHASAWNEH M., ALHAMMADI D. A., KHAN M. Z., ULLAH I., ALHARBI M.: Alzheimer's disease prediction using 3d-cnns: Intelligent processing of neuroimaging data. *SLAS Technology 32* (2025), 100265. URL: https://www.sciencedirect.com/science/article/pii/S2472630325000238, doi:https://doi.org/10.1016/j.slast.2025.100265. 1

[SCD*17] SELVARAJU R. R., COGSWELL M., DAS A., VEDANTAM R., PARIKH D., BATRA D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)* (2017), pp. 618–626. doi:10.1109/ICCV.2017.74. 2

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.