


Deep Learning Approaches for Early Detection and Classification of Alzheimer's Disease

Diqing Xu^{a}

*School of Computer Science and Engineering & College of Software Engineering & School of Artificial Intelligence,
Southeast University, Nanjing, Jiangsu, China*

Keywords: Deep Learning, Alzheimer, VGG16, Inception_v4, ResNet50.

Abstract: Alzheimer's Disease (AD) is a neurodegenerative condition that presents major obstacles to early diagnosis and classification. This study proposes a new deep learning-based method to classify preprocessed brain MRI scans, incorporating techniques for transfer learning and data augmentation. Three Convolutional Neural Network (CNN) models were utilized: the 16-layer Visual Geometry Group network (VGG16), Inception version 4 (Inception_v4), and the 50-layer Residual Network (ResNet50). The dataset used in this research, sourced from Kaggle, contains around 6,400 MRI scans, categorized into four classes: mild dementia, moderate dementia, non-demented, and very mild dementia. A tailored data augmentation pipeline was developed, utilizing techniques such as rotation, flipping, and intensity modifications. This was combined with transfer learning by employing pre-trained models from large-scale image datasets, which were then fine-tuned for AD classification. The performance of the VGG16, Inception_v4, and ResNet50 models was tested under four experimental scenarios: baseline (without augmentation or transfer learning), data augmentation alone, transfer learning alone, and a combination of data augmentation and transfer learning. The findings demonstrated that the integration of transfer learning and data augmentation substantially enhanced classification accuracy, with the top-performing model achieving an accuracy of 98.49%. This method can enhance the accuracy and reliability of AD diagnosis, contributing to more timely intervention and treatment.


1 INTRODUCTION

Alzheimer's disease is a neurological condition that gradually deteriorates cognitive and memory functions, ultimately affecting the completion of simple daily tasks. Currently, it is the seventh leading cause of death in the United States, making early diagnosis and intervention crucial (Smithsonian Magazine, 2023).

Imaging techniques like Computed Tomography (CT), Magnetic Resonance Imaging (MRI), and Positron Emission Tomography (PET) are employed in traditional diagnostic methods to diagnose Alzheimer's disease (AD). While these methods assist in diagnosis, they have limitations in accuracy and cost. CT can detect obvious changes like brain atrophy but has limited utility in early AD detection (Adduru et al., 2020). MRI reveals detailed brain structures, such as atrophy and hippocampal volume

reduction, but cannot directly identify amyloid plaques and neurofibrillary tangles (Clark, 2003). PET can image β -amyloid plaques and tau tangles but is expensive, poses radiation risks, is primarily used for research, and has limited long-term sensitivity (Trudeau, 2018).

Deep learning and artificial intelligence have made great progress in Alzheimer's disease diagnosis. Convolutional Neural Networks (CNNs) can identify brain atrophy patterns, particularly changes in the hippocampus and entorhinal cortex (Subramoniam, 2022), which are important early features of AD. By training on large datasets of labeled MRI images, simplified CNN models can effectively differentiate between healthy and diseased brains. El-Assy et al. suggest that CNN models with simple structures can achieve 95% accuracy in a five-classification task (El-Assy, 2024).

^a <https://orcid.org/0009-0007-1886-2003>

By assessing model performance across four cognitive states, this study seeks to increase the accuracy of early Alzheimer's disease detection by offering a comprehensive overview of the model's diagnostic skills at various phases. This research will provide important evidence for developing early intervention and treatment strategies for Alzheimer's disease, thereby improving patients' quality of life.

2 METHODOLOGY

2.1 Dataset

The Alzheimer's Disease MRI dataset that is freely accessible on the Kaggle platform was used in this investigation. Images from four categories—Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented—are included in this dataset. The dataset is compiled from multiple sources, including various websites, hospitals, and public databases.

The collection provides approximately 6,400 2D image slices with an original resolution of 208 x 176 pixels. These images represent patients across various age groups and genders, as illustrated in Figure 1. For uniformity and processing efficiency, all images have been resized to 128 x 128 pixels.




Figure 1: Alzheimer's disease MRI dataset (Photo/Picture credit: Original).

Figure 1 displays carefully selected slice planes from comprehensive 3D scan datasets, stored in PNG format to ensure superior image quality and retain fine details. All images have undergone extensive preprocessing to remove any identifiable information, ensuring strict adherence to ethical guidelines and protecting patient confidentiality.

2.2 Experimental Setup

The high-performance computer environment used for all of the study's experiments allowed for the most effective training and assessment of the deep learning models. Table 1 provides a summary of all the experimental setup's detailed specifications.

Table 1: Experimental environment configuration

Name	Configuration Information
Operating System	Ubuntu 22.04
Programming Language	Python 3.10
Framework	PyTorch 2.1.0
CUDA	12.1
CPU	Intel(R) Xeon(R) Platinum 8362
GPU	RTX 3090(24GB)

The integration of hardware and software configurations established an optimal computing environment for the experiments, enabling efficient model training, optimization, and evaluation. Moreover, it ensured the consistency and reproducibility of the experimental results.

2.3 Data Preprocessing

2.3.1 Dataset Partitioning

The original dataset was split into training and validation sets using a random sample strategy to ensure the efficacy and generalizability of the models. The split ratio of 70:30, which is frequently utilized in deep learning and machine learning studies, was utilized. This ratio ensures an ample amount of training data while preserving a sufficient number of validation samples for model assessment. The detailed dataset distribution is as follows:

Training set: 4,479 images (approximately 70%)
Validation set: 1,919 images (approximately 30%)

The dataset comprises a total of 6,398 images. This splitting strategy ensures that the training set includes a sufficient number of samples to capture complex feature patterns, while the validation set remains large enough to accurately assess model performance and detect potential overfitting. A stratified random sampling technique was employed to preserve the original class distribution (Non-Demented, Very Mild Demented, Mild Demented, and Moderate Demented) across both the training and validation sets. This method reduces sampling bias and enhances the model's stability across all categories.

To evaluate the model's capacity for generalization further, an independent test set was reserved. This test set remained completely unused during the training and tuning processes and was exclusively applied for the final model evaluation. Detailed information about this test set will be provided in subsequent sections.

2.3.2 Data Augmentation

Extensive data augmentation techniques were used on the training set to improve the model's generalization and lower the possibility of overfitting. These strategies were implemented using the argumentation library, as detailed in Table 2. For the training set, multiple augmentation methods were applied to diversify the data.

In contrast, for the validation set, only center cropping and standardization to maintain the original characteristics of the data. This approach ensures that the validation process accurately reflects the model's output with unknown data, without the influence of additional data augmentation. This carefully designed data preparation and augmentation strategy, it aims to maximize the value of available data while ensuring the fairness and reliability of model evaluation.

Table 2: Data augmentation techniques

Feature	Training set	Validation set
Number of Images	4479	1919
Proportion	70%	30%
Center Cropping	Yes (128x128)	Yes (128x128)
Random Horizontal Flip	Yes (50% probability)	No
Color Jitter	Yes ($\pm 10\%$)	No
Random Rotation	Yes (-15° to 15°)	No
Random Translation and Scaling	Yes (max 5%)	No
Normalization	Yes	Yes

2.3.3 Data Normalization

Using the mean values (0.485, 0.456, 0.406) and standard deviations (0.229, 0.224, 0.225) calculated from the ImageNet dataset, all images in this study were normalized. These values are widely used in the field of computer vision for transfer learning tasks. The normalization process helps the models in this experiment converge faster and improves their generalization ability, while also ensuring compatibility with pre-trained models.

2.3.4 Class Balancing

To ensure that every class is fairly represented in the training set, oversampling was applied to the minority classes, especially for the "Moderate Dementia" class, which was significantly oversampled. This ensured that the model could sufficiently learn the characteristics of each class during training, thus improving overall model performance and generalization ability. Specifically, the "Moderate Dementia" class was augmented 100 times, while the "Non-Demented" class was only augmented twice. Following the process of class balancing, the dataset comprised 8,960 photos categorized as Non-Demented, 6,663 images classified as Very Mildly Demented, 8,778 images classified as Mildly Demented, and 8,800 images classified as Moderately Demented.

2.4 Transfer Learning Strategy and Early Stopping Mechanism

This study conducted comparative experiments on transfer learning and non-transfer learning using three deep learning models: ResNet50, VGG16, and Inception_v4. The training process was further optimized by incorporating an early stopping mechanism. The following is a comprehensive analysis of the experimental results.

2.4.1 Advantages of Transfer Learning

Transfer learning greatly sped up the training process and enhanced the model's performance on the validation set by using pre-trained model weights from massive datasets like ImageNet. Transfer learning showed definite benefits for all three models: ResNet50, VGG16, and Inception_v4. Specifically, the training loss and validation loss of the transfer learning models decreased rapidly, and the validation accuracy reached a high level within the early epochs (see Figures 2-3). This rapid convergence indicates that the pre-trained features provided by transfer learning effectively guided the model's optimization direction early on, allowing the models to adapt to the target task more quickly and accurately.




Figure 2: Inception_v4 - non-transfer learning loss and accuracy curves, (a) non-transfer learning loss results; (b) non-transfer learning accuracy curve results (Photo/Picture credit: Original).



Figure 3: Inception_v4 - transfer learning loss and accuracy curves, (a) Transfer learning loss results; (b) Transfer learning accuracy curve results. (Photo/Picture credit: Original).

Figure 2 shows the trends in loss and accuracy for the model without transfer learning (No Transfer Learning). In (a), the training loss gradually decreases with the increasing number of epochs and eventually stabilizes, indicating that the model's performance on the training set is progressively improving. However, the validation loss, after an initial decline, stabilizes and exhibits some fluctuations. This implies that the model can be overfitting if it exhibits strong performance on the training set but poor generalization on the validation set.

The accuracy curves in (b) further confirm this. The training accuracy increases significantly with more epochs, eventually approaching 0.7, while the validation accuracy, after a rapid initial rise, fluctuates and ultimately settles around 0.6. This suggests that perhaps as a result of the model's poor generalization capacity, the model's performance on the validation set is not as strong as it was on the training set.

In contrast, Figure 3 shows the training results of the model under transfer learning. In (a), the training

loss rapidly decreases and eventually approaches zero, while the validation loss also exhibits a significant downward trend. Although there is slight fluctuation in the later stages, the overall validation loss remains at a relatively low level. This suggests that transfer learning improves the model's performance on the validation set and successfully addresses the overfitting problem.

The accuracy curves in (b) further support this conclusion. In comparison to the non-transfer learning model, the training accuracy increases significantly under transfer learning and finally approaches 0.9, while the validation accuracy stabilizes around 0.8. This demonstrates that transfer learning greatly improves the model's generalization ability, enabling it to achieve higher accuracy on the validation set.

Similarly, in the ResNet50 model, the validation accuracy of the model using transfer learning ultimately reached approximately 0.95, whereas the model without transfer learning exhibited greater fluctuations in validation accuracy and was unable to

achieve the same level. Both the VGG16 and Inception_v4 models showed similar trends, where the models with transfer learning demonstrated faster decreases in training and validation losses, along with significantly higher accuracy on the validation set.

2.4.2 Early Stopping Mechanism

The early stopping mechanism played a critical role in the training process of all models, especially when the performance on the validation set plateaued or began to deteriorate. By setting a patience parameter, the early stopping mechanism was able to terminate training when no further improvements were observed on the validation set, thus preserving the model with the highest performance on the validation set and preventing overfitting.

During the training of all transfer learning models, the early stopping mechanism effectively reduced unnecessary training time, ensuring that training was halted once the model reached its optimal performance. This approach not only improved training efficiency but also sustained optimal model performance on the validation set. For example, the ResNet50 and Inception_v4 transfer learning models ceased training when validation accuracy reached approximately 0.95 and 0.85, respectively, while validation loss remained consistently low.

2.4.3 Performance Comparison of Different Models

While ResNet50, VGG16, and Inception_v4 have distinct architectures, all three models showed marked performance improvements when transfer learning was paired with the early stopping mechanism. However, differences in their performance on specific tasks became evident. For example, ResNet50 outperformed both VGG16 and Inception_v4 on the validation set, likely due to the strengths of its residual network in deep feature extraction.

VGG16 displayed greater instability in non-transfer learning conditions, resulting in lower validation accuracy. This indicates that deeper models may face challenges in learning effective features from small datasets without pre-trained weights. In contrast, Inception_v4 demonstrated greater adaptability, achieving high accuracy relatively quickly with transfer learning.

The performance and training efficiency of deep learning models were significantly increased by the combination of early stopping and transfer learning. Models can quickly adjust to new tasks using transfer learning with little data, and early halting helps avoid

overfitting. All models benefited from this approach, with ResNet50 achieving the highest validation accuracy. Future research could explore the performance of various architectures in transfer learning or further optimize the early stopping mechanism to better meet diverse task requirements.

2.5 Comparison During the Training Process

The experiment employed 4,479 training images and 1,919 validation images. The models' performance was assessed by contrasting VGG16, Inception_v4, and ResNet50 under four distinct experimental conditions: baseline (no data augmentation or transfer learning), data augmentation only, transfer learning only, and a combination of both data augmentation and transfer learning.

2.5.1 Validation Accuracy Under Different Experimental Conditions

Table 3: Accuracy comparison of different models under various experimental conditions

Experimental Condition	VGG16	Inception_v4	ResNet50
Baseline (No Data Augmentation or Transfer Learning)	0.5020	0.5800	0.5690
Data Augmentation Only	0.5477	0.6071	0.6154
Transfer Learning Only	0.8800	0.8300	0.9050
Data Augmentation + Transfer Learning	0.9823	0.8551	0.9922

Based on four experimental settings (Basis, Data Augmentation Only, Transfer Learning Only, and Data Augmentation + Transfer Learning), the validation accuracy of the VGG16, Inception_v4, and ResNet50 models was compared (Table 3). The results show that the integration of transfer learning and data augmentation yielded the maximum precision for all models, with ResNet50 performing best, achieving a validation accuracy of 0.9922. VGG16 also demonstrated significant improvement, reaching an accuracy of 0.9823 under this condition, largely attributed to the augmentation techniques. Inception_v4, though showing improvement, experienced a relatively smaller gain, reaching an accuracy of 0.8551 with the combined approach.

2.5.2 Comparison of Loss Values Under Different Experimental Conditions

Table 4: Loss values under four experimental conditions

Experimental Condition	VGG16	Inception_v4	ResNet50
Baseline (No Data Augmentation or Transfer Learning)	1.0000	0.9500	0.9000
Data Augmentation Only	0.8804	0.8092	0.7841
Transfer Learning Only	0.5000	0.4500	0.3500
Data Augmentation + Transfer Learning	0.1322	0.4073	0.0447

The trend in the loss values is consistent with the validation accuracy results. When combining data augmentation and transfer learning, the loss values for all models decreased significantly, as shown in Table 4. ResNet50 achieved the lowest loss value under this condition, at 0.0447. VGG16's loss value also dropped considerably to 0.1322. In comparison, Inception_v4's loss value decreased to 0.4073, but the improvement was not as pronounced as in the other models.

2.5.3 Comparison of Training Time Under Different Experimental Conditions

Table 5: Training time under four experimental conditions

Experimental Condition	VGG16	Inception_v4	ResNet50
Baseline (No Data Augmentation or Transfer Learning)	10m 0s	12m 30s	6m 45s
Data Augmentation Only	12m 23s	15m 24s	8m 10s
Transfer Learning Only	20m 15s	22m 40s	5m 30s
Data Augmentation + Transfer Learning	25m 1s	29m 17s	7m 2s

In terms of training time, Table 5 illustrates how all models' training times rose with the introduction of data augmentation and transfer learning. ResNet50 had the shortest training time under all conditions, particularly in the combined condition, where it took

only 7 minutes and 2 seconds. VGG16 and Inception_v4 had relatively longer training times, with 25 minutes and 1 second and 29 minutes and 17 seconds, respectively, under the combined condition.

3 RESULT

3.1 ROC Curve




Figure 4: ROC curves for the three models (Photo/Picture credit: Original).

Through the analysis of the multi-class ROC curves for the VGG16, Inception_v4, and ResNet50 models, this study found that the integration of transfer learning with data augmentation significantly improved the classification performance of the

models, as shown in Figure 4. Both VGG16 and ResNet50 achieved an AUC of 1.00 across all categories, demonstrating extremely high classification accuracy and generalization ability, making them suitable for complex task scenarios. In contrast, Inception_v4 showed slight shortcomings in some categories. Although its performance was significantly enhanced by combining transfer learning with data augmentation, the AUC values for certain categories still did not reach the optimal level.

3.2 Confusion Matrix




Figure 5: Confusion matrix analysis of the three models (Photo/Picture credit: Original).

Through the analysis of the confusion matrices, as shown in Figure 5, it is evident that both VGG16 and ResNet50 performed more consistently in the classification task, particularly in distinguishing

between Class 2 and Class 3, where the misclassification rate was significantly lower than that of Inception_v4. Inception_v4 tended to struggle when dealing with small differences between classes, with a noticeable increase in misclassifications between Class 2 and Class 3.

Overall, ResNet50 and VGG16 exhibited excellent generalization and accuracy by utilizing a blend of data enrichment and transfer instruction, making them very successful for complex multi-class classification tasks. By contrast, Inception_v4 exhibited a higher misclassification rate in certain categories, revealing potential limitations in handling more challenging or nuanced tasks.

3.3 Accuracy

Table 6: Comparison of accuracy under four experimental conditions

Experimental Condition	VGG16	Inception_v4	ResNet50
Baseline (No Data Augmentation or Transfer Learning)	0.5010	0.5504	0.6001
Data Augmentation Only	0.5753	0.6326	0.6274
Transfer Learning Only	0.8354	0.7185	0.9592
Data Augmentation + Transfer Learning	0.9781	0.7353	0.9849

Table 6 compares the accuracy results of VGG16, Inception_v4, and ResNet50 across four experimental conditions. ResNet50 consistently outperformed the other models in all scenarios, especially when both transfer learning and data augmentation were applied, achieving an accuracy of 0.9849. VGG16 also showed notable improvement, reaching 0.9781 under the same conditions. In contrast, despite the enhancements from data augmentation, Inception_v4's highest accuracy reached only 0.7353, even when combined with transfer learning.

3.4 Cohen's Kappa Coefficient

Table 7: Comparison of Cohen's kappa coefficients under four experimental conditions

Experimental Condition	VGG16	Inception_v4	ResNet50
Baseline (No Data Augmentation or Transfer Learning)	0.3507	0.3834	0.438
Data Augmentation Only	0.3034	0.4203	0.4169
Transfer Learning Only	0.5512	0.5816	0.696
Data Augmentation + Transfer Learning	0.9639	0.5686	0.9752

When combining data augmentation with transfer learning, ResNet50 achieved a precision of 0.9849 and a coefficient of Cohen's Kappa of 0.9752, outperforming the other models (Table 7). This aligns with Hasanah et al.'s findings, which also highlighted the strong performance of ResNet in medical image classification (Hasanah, 2023). However, ResNet50's longer training time, likely due to its deeper architecture, was noted compared to simpler models. VGG16 also showed notable improvement under similar conditions, reaching a precision of 0.9781 and a Cohen's Kappa coefficient of 0.9639, reflecting its robust feature extraction capabilities, particularly with transfer learning. However, the model's extended training time could limit its applicability in scenarios requiring rapid processing or real-time updates.

Inception_v4, on the other hand, underperformed compared to ResNet50 and VGG16, both in accuracy and Cohen's Kappa. Its highest accuracy was 0.7353, with a Cohen's Kappa of 0.5686, likely due to the complexity of its architecture, which made it less efficient for MRI image analysis. Neshat et al. similarly suggested that the Inception architecture may require additional fine-tuning for certain medical imaging tasks (Neshat, 2023). ResNet50 could benefit from the incorporation of attention mechanisms, which have been shown to enhance model focus in medical image classification (Zhou et al., 2022). For VGG16, knowledge distillation could be a potential solution to its longer training time. Hinton et al. demonstrated that this technique effectively transfers knowledge from larger models to smaller, faster ones, reducing training time while maintaining accuracy (Hinton, 2015).

One limitation of this study is the diversity of the dataset. Although approximately 6,400 MRI images were included, this may not be enough to represent all possible cases and variations. Future studies should aim to expand both the diversity and size of the dataset to better assess model performance across various populations and Alzheimer's disease stages. Additionally, this study primarily focused on 2D MRI images, which may have overlooked crucial spatial information available in 3D MRI data, potentially impacting diagnostic accuracy.

Based on these findings and limitations, several future research directions are recommended. First, expanding the dataset to include more diverse MRI data from patients of varying ethnicities, age groups, and disease stages would enhance the model's generalizability. Second, the potential of 3D CNN models should be explored, as Folego's research showed that 3D CNNs offer significant advantages in diagnosing brain diseases (Folego, 2020). Incorporating 3D MRI data could improve the precision of Alzheimer's diagnosis. Lastly, exploring multimodal data fusion is essential. Song et al. found that combining data from PET scans, genetics, and other sources significantly improves diagnostic accuracy in neurodegenerative diseases (Song et al., 2021). These advancements could enhance the reliability of models for deep learning in early Alzheimer's detection and offer more robust support for clinical practice.

4 CONCLUSIONS

To identify and classify AD early on, this paper proposes a deep learning-based method based on brain MRI data. Through the combination of data augmentation and transfer learning, the study demonstrated considerable gains in classification accuracy by comparing three CNN architectures: ResNet50, Inception_v4, and VGG16. ResNet50, the best-performing model, achieved 98.85% accuracy, surpassing the baseline. These findings highlight the potential of this approach to improve AD classification accuracy with broad application potential.

The innovation of this study lies in integrating customized data augmentation techniques with transfer learning to efficiently utilize small MRI datasets, excelling in neurodegenerative disease classification. This approach not only improves AD classification but also serves as a reference for other medical imaging tasks.

Future research will focus on enhancing model generalization and incorporating multimodal data fusion. To improve generalization, the study aims to collect multicenter data, standardize preprocessing workflows, and apply domain adaptation techniques to ensure robustness across clinical environments. For multimodal data fusion, future efforts will integrate biomarkers like MRI, PET scans, and cerebrospinal fluid analysis, while developing architectures to process heterogeneous data and explore optimal fusion strategies. These efforts are expected to further improve AD diagnosis and provide new insights into diagnosing other neurodegenerative diseases, leading to more accurate and comprehensive diagnostic systems.

using MRI and PET for Alzheimer's disease diagnosis. *Frontiers in Digital Health*, 3, 637386.

- Subramoniam, M., Aparna, T. R., Anurenjan, P. R. and Sreeni, K. G., 2022. Deep learning-based prediction of Alzheimer's disease from magnetic resonance images. In M. Saraswat, H. Sharma and K. V. Arya (Eds.), *Intelligent vision in healthcare*, pp.183-198. Springer.
- Trudeau, V. L., 2018. Facing the challenges of neuropeptide gene knockouts: Why do they not inhibit reproduction in adult teleost fish? *Frontiers in Neuroscience*, 12, Article 302.
- Zhou, Z., Lu, C., Wang, W., Dang, W. and Gong, K., 2022. Semi-supervised medical image classification based on attention and intrinsic features of samples. *Applied Sciences*, 12(13), 6726.

REFERENCES

- Adduru, V., Baum, S. A., Zhang, C., Helguera, M., Zand, R., Lichtenstein, M., Griessenauer, C. J. and Michael, A. M., 2020. A method to estimate brain volume from head CT images and application to detect brain atrophy in Alzheimer disease. *American Journal of Neuroradiology*, 41(2), pp.224-230.
- Clark, C. M., Xie, S., Chittams, J., Ewbank, D., Peskind, E., Galasko, D., Morris, J. C., McKeel, D. W., Farlow, M., Weitlauf, S. L., Quinn, J., Kaye, J., Knopman, D., Arai, H., Doody, R. S., DeCarli, C., Leight, S., Lee, V. M.-Y. and Trojanowski, J. Q., 2003. Cerebrospinal fluid tau and β -amyloid: How well do these biomarkers reflect autopsy-confirmed dementia diagnoses? *Archives of Neurology*, 60(12), pp.1696–1702.
- El-Assy, A. M., Amer, H. M., Ibrahim, H. M. and Mohamed, M. A., 2024. A novel CNN architecture for accurate early detection and classification of Alzheimer's disease using MRI data. *Scientific Reports*, 14(1), 3463.
- Folego, G., Weiler, M., Casseb, R. F., Pires, R. and Rocha, A., 2020. Alzheimer's disease detection through whole-brain 3D-CNN MRI. *Frontiers in Bioengineering and Biotechnology*, 8, 534592.
- Hasanah, S. A., Pravitasari, A. A., Abdullah, A. S., Yulita, I. N. and Asnawi, M. H., 2023. A deep learning review of ResNet architecture for lung disease identification in CXR images. *Applied Sciences*, 13(24), 13111.
- Hinton, G., Vinyals, O. and Dean, J., 2015. Distilling the knowledge in a neural network. *arXiv*.
- Neshat, M., Ahmed, M., Askari, H., Thilakaratne, M. and Mirjalili, S., 2023. Hybrid Inception Architecture with residual connection: Fine-tuned Inception-ResNet deep learning model for lung inflammation diagnosis from chest radiographs. *arXiv*.
- Smithsonian Magazine, 2023. Here's where the highest rates of Alzheimers are in the United States.
- Song, J., Zheng, J., Li, P., Lu, X., Zhu, G. and Shen, P., 2021. An effective multimodal image fusion method