



# Is a PET All You Need? A Multi-modal Study for Alzheimer's Disease Using 3D CNNs

Marla Narazani<sup>1</sup>(✉), Ignacio Sarasua<sup>1,2</sup>, Sebastian Pölsterl<sup>2</sup>,  
Aldana Lizarraga<sup>1</sup>, Igor Yakushev<sup>1</sup>, and Christian Wachinger<sup>1,2</sup>

<sup>1</sup> School of Medicine, Technical University of Munich, Munich, Germany  
[marla.narazani@tum.de](mailto:marla.narazani@tum.de)

<sup>2</sup> Lab for Artificial Intelligence in Medical Imaging (AI-Med), KJP, LMU Klinikum,  
Munich, Germany

**Abstract.** Alzheimer's Disease (AD) is the most common form of dementia and often difficult to diagnose due to the multifactorial etiology of dementia. Recent works on neuroimaging-based computer-aided diagnosis with deep neural networks (DNNs) showed that fusing structural magnetic resonance images (sMRI) and fluorodeoxyglucose positron emission tomography (FDG-PET) leads to improved accuracy in a study population of healthy controls and subjects with AD. However, this result conflicts with the established clinical knowledge that FDG-PET better captures AD-specific pathologies than sMRI. Therefore, we propose a framework for the systematic evaluation of multi-modal DNNs and critically re-evaluate single- and multi-modal DNNs based on FDG-PET and sMRI for binary healthy vs. AD, and three-way healthy/mild cognitive impairment/AD classification. Our experiments demonstrate that a single-modality network using FDG-PET performs better than MRI (accuracy 0.91 vs 0.87) and does not show improvement when combined. This conforms with the established clinical knowledge on AD biomarkers, but raises questions about the true benefit of multi-modal DNNs. We argue that future work on multi-modal fusion should systematically assess the contribution of individual modalities following our proposed evaluation framework. Finally, we encourage the community to go beyond healthy vs. AD classification and focus on differential diagnosis of dementia, where fusing multi-modal image information conforms with a clinical need.

## 1 Introduction

With life expectancies rising globally, dementia is becoming a growing concern for individuals and society. Dementia is characterized by a progressive cogni-

M. Narazani and I. Sarasua—These authors contributed equally to this work.

**Supplementary Information** The online version contains supplementary material available at [https://doi.org/10.1007/978-3-031-16431-6\\_7](https://doi.org/10.1007/978-3-031-16431-6_7).

tive impairment that eventually requires individuals to be completely dependent upon caregivers. While this process cannot be reversed, recent efforts have focused on diagnosing subjects at an early stage to improve disease management [4]. A particular focus has been on Alzheimer’s Disease (AD), given that it is the most common form of dementia and benefits from large data-sharing initiatives [21]. To date, a wide range of diagnostic tools are available for diagnosing AD: magnetic resonance imaging (MRI), positron emission tomography (PET), cerebrospinal fluid (CSF), demographics, cognitive tests, and genetic alterations [1]. Structural MRI (sMRI) captures regional atrophy of the brain, whereas FDG-PET measures the brain’s glucose metabolism. FDG-PET plays a major role in the clinical diagnosis of AD. It can detect functional brain changes in AD early in the disease progression and can help to differentiate AD from other causes of dementia such as frontotemporal and Lewy body dementia [22]. In the memory clinic, MRI and FDG-PET are among the most common neuroimaging methods used [26], where FDG-PET is considered to have a higher diagnostic and prognostic accuracy [3, 12].

Recently, studies on deep learning (DL) techniques have emerged that showed that distinguishing healthy controls from AD subjects becomes more accurate when learning from MRI *and* FDG-PET, rather than a single modality [24, 29, 30]. However, this scenario is very different from that in a memory clinic. In the clinic, the main objective is differential diagnosis to determine the type of dementia, whereas studies on DL merely considered a single type of dementia, namely AD [24, 29, 30]. When considering that both modalities assess neural degeneration, but AD-specific changes are better captured by FDG-PET than MRI [3, 12], it seems surprising why combining MRI and FDG-PET with DL would be beneficial when AD is the only form of dementia that is being studied.

In this work, we critically re-evaluate single- and multi-modal DL models based on FDG-PET and structural MRI for classifying healthy vs. AD subjects. We study three different modes of multimodal fusion: early, middle, and late fusion. We evaluate each to investigate whether it truly benefits from multimodal data by performing ablation studies for which MRI and FDG-PET images are paired randomly. Contrary to previous work, our experiments show that FDG-PET alone is sufficient for AD diagnosis, which conforms with established clinical knowledge about biomarkers in AD. We argue that future work on multimodal fusion should follow our proposed evaluation framework to systematically assess the contribution of individual modalities.

**Related Work.** Most DL models for AD prediction are single-modal (see [7] for an overview). In [9], the authors propose a 2D convolutional neural network (CNN) using slices of sMRI volumes. However, recent work has shifted towards 3D CNN architectures for analyzing sMRI [2, 8, 14, 18, 19, 23]. A sparse autoencoder is combined with a CNN in [23]. Korolev et al. [18] compare a 3D-VGG and 3D-Resnet architecture. Both [2] and [8] use a 3D CNN for whole brain MRIs. Regarding work related to FDG-PET, a 2D CNN has been used in [6, 20], and a 3D CNN in [28]. Finally, several works combined sMRI and FDG-PET

[24, 29, 30]. In [24], the authors propose an early fusion approach by overlaying gray matter (GM) tissues from MRI with the FDG-PET scans and evaluate the effectiveness of their fusion strategy using a 3D CNN. In [30], a three-stage framework based on middle and late fusion using MRI, FDG-PET, and single nucleotide polymorphisms is proposed. The authors of [10] combine a 3D CNN and LSTM. Finally, in [15], an early and a late fusion approach are presented based on a 3D-VGG. The works on multi-modal fusion unanimously concluded that fusing sMRI and FDG-PET improves prediction accuracy over using a single modality, which conflicts with the established clinical knowledge that FDG-PET better captures AD-specific pathologies than sMRI [3, 12].

## 2 Methods

To determine the contribution of each modality in a multi-modal DNN, we propose a systematic evaluation framework. First, we consider each modality in isolation by using a single branch 3D CNN. Next, we consider the joint contribution of multiple modalities using a 3D CNN with either early, late, or middle fusion (see Fig. 1). To assess whether multi-modal inputs are truly helpful, we perform ablation experiments where MRI and FDG-PET images are paired randomly. This allows us to quantify the importance of each modality.

### 2.1 CNN Architecture

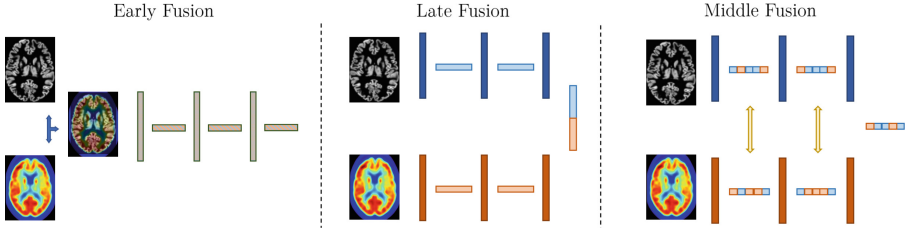
We use a 3D ResNet as the base architecture for all models (more details in supplemental Fig. S1). It comprises 12 convolutional layers with kernel size  $3^3$  in total. We use four residual learning blocks consisting of two convolutional layers followed by batch normalization (BN) [16] and rectified linear unit (ReLU) activation. We half the spatial resolution of feature maps in the last three residual blocks by using a stride of 2. Finally, we perform global average pooling across the spatial dimensions of the feature maps and use two linear layers to output a log-probability. We use dropout in each residual block to reduce overfitting.

### 2.2 Fusion Strategies

We consider three strategies for fusing multi-modal data: early, late, and middle fusion (see Fig. 1). All three strategies follow the base CNN architecture described above. Next, we describe the fusion strategies in detail.

**Early Fusion.** In early fusion, raw modalities are combined directly before being passed to the network. Here, we follow the strategy proposed in [24]: gray matter maps are obtained via Voxel-Based Morphometry (VBM) and used to mask the FDG-PET intensities. In the resulting volume, the intensities of the FDG-PET are effectively weighted by the MRI intensities. The network is a single branch network that receives the combined MRI-FDG-PET volume as input.

**Late Fusion.** Late fusion is the most straight-forward approach to fuse multi-modal data. Rather than fusing the images, it fuses the latent representations of



**Fig. 1.** Overview of the three fusion strategies. Early fusion combines the MRI and PET inputs in a single volume. Late fusion concatenates the latent representations coming from each independent network. Middle fusion exchanges channels of the intermediate feature maps along the network.

two separate networks. Here, we train two independent 3D ResNet branches, one for MRI and one for FDG-PET. The features obtained from each branch after global average pooling are then concatenated and passed through a Multi-layer perceptron (MLP) [128, 64, number of classes] to obtain a log-probability that accounts for both sources of information.

**Middle Fusion.** While early and late fusion are common in multi-modal analysis, we also explore an approach that fuses intermediate representations of modality-specific networks [27]. In this approach, modality-specific information are fused by dynamically exchanging feature maps between sub-networks of different modalities. This bi-directional exchange of information is self-guided by considering individual channel importance, which is measured by the magnitude of the BN scaling factor. This process is carried out under the  $\ell_1$  regularization to penalize exchanging all channels. To further encourage sharing of information, convolutional filter weights are shared across modalities. Note that BN layers are not shared in order to determine the channel importance for each individual modality. To the best of our knowledge, channel exchange has not been applied for multi-modal fusion for AD prediction before.

### 2.3 Evaluation Scheme

Our main objective is to rigorously evaluate whether MRI is truly relevant for diagnosing AD when FDG-PET is available too. For all of our experiments, we train the models using FDG-PET and MRI data from the same patient. During inference we define three different experiments based on the input data: (i) correct, (ii) random PET, and (iii) random MRI. We use balanced accuracy (BACC) to assess the predictive performance of models, because it is insensitive to the relative frequency of class labels [5].

**Correct.** This strategy follows the standard training and testing scheme. For each fusion strategy, we test the networks based on FDG-PET and MRI scans from the same patient. If both modalities would be relevant for AD diagnosis, we would expect this scenario to yield the highest predictive performance. It serves as a baseline for the remaining experiments.

**Random MRI.** In this experiment, we pair a patient’s true FDG-PET image and diagnosis with an MRI of a randomly selected patient. If both modalities would be relevant for the final decision, we would expect a significant drop in performance with respect to the *Correct MRI* experiment. Otherwise, if performance remains similar, the contribution of patient-specific anatomy, as captured by the MRI, the MRI adds little additional information that is not available from the FDG-PET.

**Random PET.** This experiment is similar to the previous experiment, but this time we pair the correct MRI and diagnosis with a randomly selected FDG-PET from another patient. The conclusions we can derive from this experiment are the same as in the previous experiment, but focus on assessing the contribution of FDG-PET.

## 2.4 Data Processing and Training Strategy

We use pre-processed FDG-PET scans and T1-weighted MRI scans from the Alzheimer’s disease neuroimaging initiative (ADNI; [17]) database. Full details about the pre-processing steps can be found at the ADNI website for FDG-PET<sup>1</sup> and for MRI<sup>2</sup>. Both scans were additionally processed using SPM<sup>3</sup> and CAT12 [13]. MRI scans were processed using the standard VBM pipeline in CAT12<sup>4</sup>. We use the gray matter (GM) tissue area of the brain as an input to the network. FDG-PET scans were normalized and registered to the MNI152 template [11] with 1.5 mm<sup>3</sup> voxel size. We performed min-max scaling to rescale the image intensity values to the range between 0 and 1. The final image size for both modalities is 113 × 137 × 113.

Our dataset comprises 257 patients with AD, 370 healthy controls (CN), and 611 patients with mild cognitive impairment (MCI); see the supplemental Table S1 for additional information. We split the data into train/validation/test sets with sizes roughly in 65/15/20% of the full dataset. We perform cross-validation across 5 splits, based on a data stratification strategy that accounts for sex, age and diagnosis. We only include baseline visits scans so that only a single scan per patient is available. We train models for two tasks (i) binary classification of healthy controls (CN) vs. patients with AD, and (ii) three-way classification of CN vs. MCI vs. AD. All models are trained end-to-end using a cross-entropy loss and data augmentation during training (up to 8° angle rotation and 8 mm translation in each dimension). More information about the training setup can be found in the supplemental Table S2.

## 3 Results

**Testing on Random PET or MRI.** Table 1 reports the results for the experiments described in Sect. 2.3, for binary and three-way classification. We observe

<sup>1</sup> <http://adni.loni.usc.edu/methods/pet-analysis>.

<sup>2</sup> <http://adni.loni.usc.edu/methods/mri-analysis>.

<sup>3</sup> <https://www.fil.ion.ucl.ac.uk/spm/software/spm12>.

<sup>4</sup> <http://www.neuro.uni-jena.de/cat12/CAT12-Manual.pdf>.

**Table 1.** Overview of the evaluation scheme for correct data, random MRI or random PET. Numbers are mean balanced accuracy (BACC) and standard deviation across folds.

	Random MRI	Random PET	BACC 2-Class	BACC 3-Class
Early fusion	✗	✗	$0.885 \pm 0.041$	$0.573 \pm 0.023$
Early fusion	✗	✓	$0.696 \pm 0.026$	$0.414 \pm 0.015$
Early fusion	✓	✗	$0.720 \pm 0.015$	$0.470 \pm 0.028$
Middle fusion	✗	✗	$0.893 \pm 0.036$	$0.530 \pm 0.034$
Middle fusion	✗	✓	$0.527 \pm 0.020$	$0.366 \pm 0.028$
Middle fusion	✓	✗	$0.890 \pm 0.020$	$0.528 \pm 0.025$
Late fusion	✗	✗	$0.896 \pm 0.019$	$0.577 \pm 0.029$
Late fusion	✗	✓	$0.597 \pm 0.029$	$0.368 \pm 0.027$
Late fusion	✓	✗	$0.786 \pm 0.080$	$0.527 \pm 0.038$

that when testing on the correct pair of scans, all fusion approaches perform similarly for both tasks with two exceptions: Early Fusion achieves a mean BACC approximately 0.01 lower for binary classification, and Middle Fusion a BACC approximately 0.04 lower for three-class classification. Overall, we observe a significant drop in performance between these two tasks, which is expected given that MCI is not a true diagnosis, but a syndrome, which makes it highly heterogeneous, especially with limited amount of training data.

Interestingly, if we look at the results for the middle and late fusion models when testing on partially random data, we observe a much larger drop in performance when the FDG-PET is randomized; the accuracy is close to random chance. On the other hand, randomizing the MRI data has much lower impact on the overall performance. For binary classification the mean BACC drops around 0.11 for late fusion and merely 0.003 for middle fusion, which is much lower than for the random PET experiments: 0.299 and 0.366, respectively. For early fusion, results for both randomized experiments experience a significant drop compared to using the original data. This outcome is expected, since early fusion results in a single volume where the MRI acts as a mask to select regions from the FDG-PET. If the pair of images is from different patients, anatomies are not perfectly aligned and early fusion will remove important areas. Hence, the effect of randomizing the MRI or the FDG-PET leads to a similar loss in information and comparable drop in performance.

**Training on Random MRI:** The performance difference between randomizing the FDG-PET data vs. the MRI (see Table 1) suggests that both modalities do not have the same contribution to the models’ final decision. We decided to further evaluate this hypothesis by defining an additional experiment: during training, the FDG-PET remains associated to a specific patient, but the MRI is exchanged with a random subject. Table 2 shows the results for two- and three-class. Note that results for the original data (Correct) are identical to those

in Table 1. For binary classification with correct data, middle and late fusion outperform early fusion by at least 0.08 in mean BACC. Single modality PET yields the best performance on correct data. When using a random MRI, the BACC for early fusion decreases, but improves for late fusion, matching the BACC of the single modality PET. For three-classes with correct data, using PET and MRI data performs similarly with a 0.03 improvement for early and late fusion, while middle fusion decreases in performance by 0.01 compared to using only PET. For random MRI, we observe a strong improvement for middle fusion (0.168) and late fusion (0.081), while the accuracy for early fusion decreases to chance level.

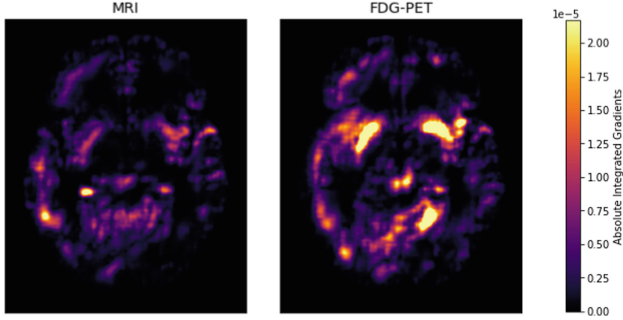
**Table 2.** Training and testing on correct, and random MRI. Numbers are mean balanced accuracy and standard deviation across folds.

	CN vs. AD		CN vs. MCI vs. AD	
	Correct	Random MRI	Correct	Random MRI
PET only	$0.905 \pm 0.015$	—	$0.541 \pm 0.034$	—
MRI only	$0.866 \pm 0.029$	—	$0.536 \pm 0.062$	—
Early fusion	$0.885 \pm 0.041$	$0.729 \pm 0.034$	$0.573 \pm 0.023$	$0.365 \pm 0.037$
Middle fusion	$0.893 \pm 0.036$	$0.863 \pm 0.026$	$0.530 \pm 0.034$	$0.698 \pm 0.087$
Late fusion	$0.896 \pm 0.019$	$0.906 \pm 0.022$	$0.577 \pm 0.029$	$0.658 \pm 0.015$

**Post-hoc Explanation via Relevance Maps:** Relevance maps are a helpful way of assessing the decision-making process of a classification model. In this work, we use them to quantify how much individual modalities contribute to the final prediction of the network. We use Integrated Gradients (IG; [25]) because its axiomatic approach allows us to precisely quantify how much the MRI and FDG-PET of a multi-modal CNN contribute to a particular prediction. Given a patient’s images and a baseline, which is defined by the user (in our case a black volume), IG computes voxel-wise contributions by integrating along the path from the baseline input to the real input. Since the sum of all voxel-wise IG scores equals the predicted log-probability, we can summarize the total contribution of the MRI and FDG-PET by summing over the IG scores for the respective modality. Figure 2 depicts the average absolute importance across 42 correctly classified AD patients by the late fusion model for CN vs. AD. This example clearly illustrates that the PET contributes significantly more to the overall predictions. Overall, the PET contributes 1.77 times more to a prediction than the MRI (sum of  $-IG-$  is 33.8 vs. 19.1), which confirms our results from above.

## 4 Discussion

We performed a thorough evaluation of the different methods across 5 splits. In our first set of experiments, we observed that when training on correct data but



**Fig. 2.** Mean absolute integrated gradients across 42 correctly classified AD patients by the late fusion model. Illustrated is an axial slice located at the center of the volume.

introducing random FDG-PET or MRI data at test time, both the middle and late networks were more sensitive to changes of PET. While this is already a strong indicator of the bias of the neural network, our second set of experiments (Table 2) give us better insights on the reasons behind this phenomenon. First of all, the AD vs. CN classification experiments are consistent across Tables 1 and 2, which proves that the middle and late fusion networks rely mostly on FDG-PET. These results are supported by the relevance maps in Fig. 2.

For the three-class experiment, the BACC is below 60%, confirming the difficulty of the task. MCI subjects are a heterogeneous group that may also suffer from other types of dementia. Therefore, the amount of data required to train a predictive model for this task is much larger than in a two-class setting. For this challenging task, the usage of random MRIs led to a steep increase in accuracy for middle and late fusion. We believe that randomizing the MRI data serves as an augmentation mechanism during training. Given that in each epoch, the model sees a different pair of FDG-PET and MRI scans, this is likely making the networks more robust to alterations during inference.

Our results, while being aligned to previous medical findings, are in disagreement with previous literature that favored the fusion of MRI and FDG-PET for AD prediction. One reason for this difference could be that randomly exchanging image pairs during training leads to a larger effective training data size, which in turn allows the network to be more robust to changes in the data distribution during inference (similarly observed in Table 2 for the random MRI experiments). Additionally, by increasing the number of branches (e.g. two branches for the late fusion) the amount of trainable parameters is almost doubled, which allows the network to define more complex decision boundaries. This also makes the networks more prone to overfitting as observed in the three-class experiment when comparing late fusion on correct or random MRI. Finally, another potential reason is the importance of the PET pre-processing. For instance, [24] use a different pre-processing for the image fusion (for which they report high accuracy) and different input for the uni-modal and concatenation networks. GM is



used in image fusion and MNI-MRI for the other combinations. When we compared the performance between processed and un-processed PET data, we noted a decrease of about 7% (t-test  $P=0.01$ ) in balanced accuracy.

## 5 Conclusion

In this work, we rigorously evaluated single- and multi-modal deep neural networks for AD diagnosis based on MRI and FDG-PET. Our results demonstrate that a single-modality network using FDG-PET performs best for healthy/AD classification. While this is in contrast with previous work on deep learning for modality fusion, it does conform with the established clinical knowledge that FDG-PET better captures AD-specific patterns of neurodegeneration than MRI. We argue that recent work on multi-modal fusion, while technically sound, are largely disconnected from the established clinical knowledge about biomarkers in AD. We argue that future work on multi-modal fusion for AD diagnosis should take the existing clinical knowledge better into account, and systematically assess the contribution of individual modalities following our experimental setup. In the future, we plan to conduct experiments for MCI vs. NC, validate our hypotheses on different datasets and test other classification models.

**Acknowledgment.** This research was partially supported by the Bavarian State Ministry of Science and the Arts and coordinated by the bidt, and the Federal Ministry of Education and Research in the call for Computational Life Sciences (DeepMentia, 031L0200A). We gratefully acknowledge the computational resources provided by the Leibniz Supercomputing Centre ([www.lrz.de](http://www.lrz.de)).

## References

1. Aisen, P.S., Cummings, J., Jack, C.R., Morris, J.C., Sperling, R.: On the path to 2025: understanding the Alzheimer’s disease continuum. *Alzheimers Res. Ther.* **9**(1), 60 (2017)
2. Basaia, S., et al.: Automated classification of Alzheimer’s disease and mild cognitive impairment using a single MRI and deep neural networks. *NeuroImage Clin.* **21**, 101645 (2019)
3. Bloudek, L.M., Spackman, D.E., Blankenburg, M., Sullivan, S.D.: Review and meta-analysis of biomarkers and diagnostic imaging in Alzheimer’s disease. *J. Alzheimer’s Dis.* **26**(4), 627–645 (2011)
4. Borson, S., et al.: Improving dementia care: the role of screening and detection of cognitive impairment. *Alzheimer’s Dement* **9**(2), 151–159 (2013)
5. Brodersen, K.H., Ong, C.S., Stephan, K.E., Buhmann, J.M.: The balanced accuracy and its posterior distribution. In: 20th International Conference on Pattern Recognition, pp. 3121–3124 (2010)
6. Ding, Y., Sohn, J.H., Kawczynski, M.G., Trivedi, H., Harnish, R., et al.: A deep learning model to predict a diagnosis of Alzheimer disease by using 18F-FDG PET of the brain. *Radiology* **290**(2), 456–464 (2019)

7. Ebrahimighahnavieh, M.A., Luo, S., Chiong, R.: Deep learning to detect Alzheimer's disease from neuroimaging: a systematic literature review. *Comput. Methods Programs Biomed.* **187**, 105242 (2020)
8. Esmailzadeh, S., Belivanis, D.I., Pohl, K.M., Adeli, E.: End-to-end Alzheimer's disease diagnosis and biomarker identification. In: *MLMI*, pp. 337–345 (2018)
9. Farooq, A., Anwar, S., Awais, M., Rehman, S.: A deep CNN based multi-class classification of Alzheimer's disease using MRI. In: *IST*, pp. 1–6 (2017)
10. Feng, C., et al.: Deep learning framework for Alzheimer's disease diagnosis via 3D-CNN and FSBi-LSTM. *IEEE Access* **7**, 63605–63618 (2019)
11. Fonov, V., et al.: Unbiased average age-appropriate atlases for pediatric studies. *Neuroimage* **54**(1), 313–327 (2011)
12. Frisoni, G.B., et al.: Imaging markers for Alzheimer disease: which vs how. *Neurology* **81**(5), 487–500 (2013)
13. Gaser, C., Dahnke, R., et al.: Cat-a computational anatomy toolbox for the analysis of structural MRI data. *HBM* **2016**, 336–348 (2016)
14. Hosseini-Asl, E., Gimel'farb, G., El-Baz, A.: Alzheimer's disease diagnostics by a deeply supervised adaptable 3D convolutional network. *arXiv preprint arXiv:1607.00556* (2016)
15. Huang, Y., et al.: Diagnosis of Alzheimer's disease via multi-modality 3D convolutional neural network. *Front. Neurosci.* **13**, 509 (2019)
16. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: *ICML*, pp. 448–456 (2015)
17. Jack, C.R., Jr., et al.: The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. *J. Magn. Resonan. Imaging* **27**(4), 685–691 (2008)
18. Korolev, S., Safullin, A., Belyaev, M., Dodonova, Y.: Residual and plain convolutional neural networks for 3D brain MRI classification. In: *ISBI*, pp. 835–838 (2017)
19. Li, F., Cheng, D., Liu, M.: Alzheimer's disease classification based on combination of multi-model convolutional networks. In: *IST*, pp. 1–5 (2017)
20. Liu, M., Cheng, D., Yan, W., et al.: Classification of Alzheimer's disease by combination of convolutional and recurrent neural networks using FDG-PET images. *Front. Neuroinform.* **12**, 35 (2018)
21. Livingston, G., Sommerlad, A., Orgeta, V., Costafreda, S.G., Huntley, J., et al.: Dementia prevention, intervention, and care. *The Lancet* **390**(10113), 2673–2734 (2017)
22. Marcus, C., Mena, E., Subramaniam, R.M.: Brain PET in the diagnosis of Alzheimer's disease. *Clin. Nucl. Med.* **39**(10), e413 (2014)
23. Payan, A., Montana, G.: Predicting Alzheimer's disease: a neuroimaging study with 3D convolutional neural networks. *arXiv preprint arXiv:1502.02506* (2015)
24. Song, J., Zheng, J., Li, P., Lu, X., Zhu, G., Shen, P.: An effective multimodal image fusion method using MRI and PET for Alzheimer's disease diagnosis. *Front. Digit Health* **3**, 19 (2021)
25. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: *ICML*, pp. 3319–3328 (2017)
26. Teipel, S., Kilimann, I., Thyrian, J.R., Kloppel, S., Hoffmann, W.: Potential role of neuroimaging markers for early diagnosis of dementia in primary care. *Curr. Alzheimer Res.* **15**(1), 18–27 (2017)
27. Wang, Y., Huang, W., Sun, F., Xu, T., Rong, Y., Huang, J.: Deep multimodal fusion by channel exchanging. *NeurIPS* **33**, 4835–4845 (2020)
28. Yee, E., et al.: Quantifying brain metabolism from FDG-PET images into a probability of Alzheimer's dementia score. *Hum. Brain Mapp.* **41**(1), 5–16 (2020)

29. Zhang, D., et al.: Multimodal classification of Alzheimer's disease and mild cognitive impairment. *Neuroimage* **55**(3), 856–867 (2011)
30. Zhou, T., Thung, K.H., Zhu, X., Shen, D.: Effective feature learning and fusion of multimodality data using stage-wise deep neural network for dementia diagnosis. *Hum. Brain Mapp.* **40**(3), 1001–1016 (2019)