Deep Learning for Diagnosis of Alzheimer's Disease with FDG-PET Neuroimaging

José Bastos, Filipe Silva, and Petia Georgieva(✉)

University of Aveiro, Aveiro, Portugal
{bastosjose,fmsilva,petia}@ua.pt
http://www.ua.pt

Abstract. Alzheimer's Disease (AD) imposes a heavy burden on health services both due to the large number of people affected as well as the high costs of medical care. Recent research efforts have been dedicated to the development of computational tools to support medical doctors in the early diagnosis of AD. This paper is focused into studying the capacity of Deep Learning (DL) techniques to automatically identify AD based on PET neuroimaging. PET images of the cerebral metabolism of glucose with fluorodeoxyglucose (^{18}F-FGD) were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database. Two DL approaches are compared: a 2D Inception V3 pre-trained model and a custom end-to-end trained 3D CNN to take advantage of the spatial patterns of the full FDG-PET volumes. The results achieved demonstrate that the PET imaging modality is suitable indeed to detect early symptoms of AD. Further to that, the carefully tuned custom 3D CNN model brings computational advantages, while keeping the same discrimination capacity as the exhaustively pre-trained 2D Inception V3 model.

Keywords: Alzheimer's disease · FDG-PET neuroimaging · Convolutional Neural Networks · ADNI dataset

1 Introduction

Neurodegenerative diseases are a spectrum of brain disorders that cause a progressive loss of neurological function and structure, such as Alzheimer's disease (AD), Parkinson's disease (PD), Huntington's disease (HD) and amyotrophic lateral sclerosis (ALS). Amongst them, AD is documented as the most common cause of dementia worldwide (responsible for 60 to 80% of cases), affecting roughly 30% of people over the age of 85 [1]. Dementia refers to a set of symptoms marked by decline in memory, reasoning or other cognitive functions. Nowadays, there is a broad consensus that AD appears decades before its first manifestation.

This research work is funded by National Funds through the FCT - Foundation for Science and Technology, in the context of the project UIDB/00127/2020.

A. J. Pinho et al. (Eds.): IbPRIA 2022, LNCS 13256, pp. 95–107, 2022.
https://doi.org/10.1007/978-3-031-04881-4_8

Apart from the search for a cure, the most recent efforts are aimed at developing computational tools to support the medical decision. Over the last few years, deep learning (DL) - based methods have made important contributions in medical imaging. They proved to be a valuable technology to assist the preventive healthcare with computerized diagnosis. In this context, Magnetic Resonance Imaging (MRI) and Positron Emission Tomography (PET) are the most common neuroimaging modalities useful for the AD diagnosis.

Several studies have highlighted the importance of DL-based diagnostic systems using either MRI or PET scans [10]. Others address the integration of multimodal information, such as PET and T1-weighted MRI images [12]. Recently, ^{18}F-FDG PET revealed to have a potential to assess the risk of AD at a very early stage [9]. PET images of the cerebral metabolism of glucose with ^{18}F-FDG provide representations of neuronal activity closely linked to the initial manifestations of AD [20].

These recent findings motivated the present work aiming to explore the potential of DL techniques in the diagnosis of AD with ^{18}F-FDG PET images. The study focuses on how to leverage convolutional neural networks (CNNs) for classifying healthy versus AD patients, with a limited dataset collected from the ADNI. For this purpose, two CNN models are compared in terms of predictive performance. The first CNN model explores transfer learning as a promising solution to the data challenge using a pre-trained model. The second model involves a custom developed 3D-CNN to take advantage of spatial patterns on the full PET volumes by using 3D filters and 3D pooling layers.

The rest of the paper is organised as follows. Section 2 reviews related works. Section 3 explains the proposed CNN framework. The results using the ADNI dataset are presented and discussed in Sect. 4. Section 5 summarises the work.

2 Related Work

The current diagnosis of AD relies on neuropsychological tests and neuroimaging biomarkers. The diagnostics can be performed in an early stage, even in the prodromal stage of the disease also referred to as mild cognitive impairment (MCI). The biomarkers for early AD diagnosis that are currently in use reflect the deposition of amyloid (CSF Aβ1-42 or PET with amyloid ligands), formation of neurofibrillary tangles (CSF P-tau), neuronal degeneration (CSF T-tau), changes in brain metabolism (FDG-PET), as well as volumetric changes in brain structures that cause the disease's symptoms, such as the hippocampus.

PET is an imaging modality that involves the application of a radioactive substance, called a radioactive tracer, into the body and the posterior observation of the emitted radiation in the organ or tissue being examined. Fluorine-18 is radioactive tracer commonly attached to compounds like glucose, as is the case with ^{18}F-FDG, for the measurement of brain metabolism.

Decreased brain glucose consumption, known as hypometabolism, is seen as one of the earliest signs of neural degeneration, being associated with AD progression. FDG-PET represents a valuable and unique tool able to estimate local

cerebral rate of glucose consumption. Thereby, PET may point out biochemical changes that underlie the onset of a disease before anatomical changes can be detected by other modalities such as CT or MRI.

Table 1. CNN applications in brain medical imaging.

Task	Modality	Reference
Tissue necrosis after CVA prediction	MRI	Stier et al. [26]
PD identification	SPECT	Choi et al. [6]
Brain tumor segmentation	MRI	Havaei et al. [11]
Brain lesion segmentation	MRI	Kamnitsas et al. [13]
Brain age prediction	MRI	Cole et al. [8]

Note: CVA = Cerebrovascular Accident
MRI = Magnetic Resonance Imaging, PD = Parkinson's Disease
SPECT = Single-Photon Emission Computed Tomography

CNNs are an important tool in medical imaging [5] and disease diagnostics, as shortly summarised in Table 1. A systematic review of deep learning techniques for the automatic detection of AD can be found in [10]. Authors emphasize important aspects to understand the whole scenario of AD diagnosis. First, approximately 73% of neuroimaging studies have been performed with single-modality data, around 83% of the studies are based on MRI, 9% refer to fMRI and only 8% to PET scans.

Second, a significant part of the studies, summarised in Table 2, address the binary classification problem, i.e., consider normal cognitive (NC) state against AD. A more challenging task is to discriminate between early and late stages of mild cognitive impairment (MCI). MCI is sometimes subdivided into sMCI (Stable Mild Cognitive Impairment) and pMCI (Progressive Mild Cognitive Impairment) which will eventually develop into AD.

Third, a common approach is to convert the volumetric data into a 2D image to be applied at the input of a 2D-CNN. Most of the studies transfer the weights from pre-trained networks on the ImageNet database to the target medical task. This process, known as transfer learning, speeds up training and reduces costs by leveraging previous knowledge.

Although the application of DL techniques in AD diagnosis is still in their initial stage, recent works [9,14] demonstrate that deep neural networks can outperform radiologist abilities. The coming years may determine the feasibility of these models as a support tool to help clinicians reach an appropriate decision in real clinical environments.

3 Methodology

As illustrated in Table 2, MRI is the most frequently used imaging modality as well as a 2D CNN as the discrimination model. In contrast, in this work we aim to leverage the full information by exploring 3D CNN for learning representations from the less explored FDG-PET data. For that purpose, a comparative study will be carried out centred on two CNN models: fine tuning of a pre-trained 2D-CNN model against an end-to-end trained from scratch custom 3D-CNN.

Table 2. CNN application for diagnosis of Alzheimer's Disease (AD).

Modality	Classes	Score	Reference
MRI	MCI vs CN	83%[a]	Qiu et al. [22]
MRI	Multi-class[b]	57%	Valliani et al. [27]
MRI	AD vs CN	91%	Liu et al. [19]
MRI	MCI vs CN	74%[a]	Li et al. [15]
MRI	sMCI vs pMCI	74%	Liu et al. [18]
MRI	sMCI vs pMCI	80%[a]	Lian et al. [16]
MRI	AD vs CN	91%	Aderghal et al. [2]
MRI	AD vs MCI	70%	Aderghal et al. [2]
MRI	MCI vs CN	66%	Aderghal et al. [2]
MRI	sMCI vs pMCI	73%	Lin et al. [17]
MRI	AD vs CN	90%	Bäckström et al. [3]
MRI	AD vs MCI	76%	Senanayake et al. [23]
MRI	MCI vs CN	75%	Senanayake et al. [23]
MRI	sMCI vs pMCI	62%	Shmulev et al. [24]
AV-45 PET	AD vs CN	85%	Punjabi et al. [21]
AV-45 PET + MRI	AD vs CN	92%	Punjabi et al. [21]
AV-45 + FDG PET	sMCI vs pMCI	84%	Choi et al. [7]

Note: CN = Cognitively Normal, AD = Alzheimer's Disease
MCI = Mild Cognitive Impairment, sMCI = Stable MCI, pMCI = Progressive MCI
[a] = Severely Imbalanced Dataset, [b] = AD vs MCI vs CN

3.1 2D Slice-Level CNN Model

The first approach is implemented by the Google 2D Inception V3 model, pre-trained on the ImageNet dataset and fine-tuned with the ADNI dataset. This approach requires a pre-processing step in which the 3D PET volume is converted into a 2D image which is the input to the pre-trained model. Inspired by the work of Ding et al. [9], a 2D collage of a grid of 4 by 4 FDG-PET scan slices were used as the inputs to a deep model, as shown in Fig. 1.

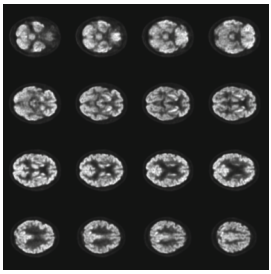

Fig. 1. Collage of 2D slices extracted from volumetric PET scans.

The advantage of this approach is that pre-trained models exist and they can be quickly updated to fit new target data [4]. Further to that, the training data is increased since larger number of 2D slices can be obtained from a single 3D sample.

3.2 3D Subject-Level CNN Model

The second approach is a custom 3D-CNN to take advantage of the spatial patterns of the full PET volumes for each subject. It is referred here as the 3D subject-level CNN model (see Fig. 2). In contrast to the 2D Slice -level model, where the raw PET scans were transformed into 2D patches, here the data is first pre-processed into 3D tensors and then loaded into the network.

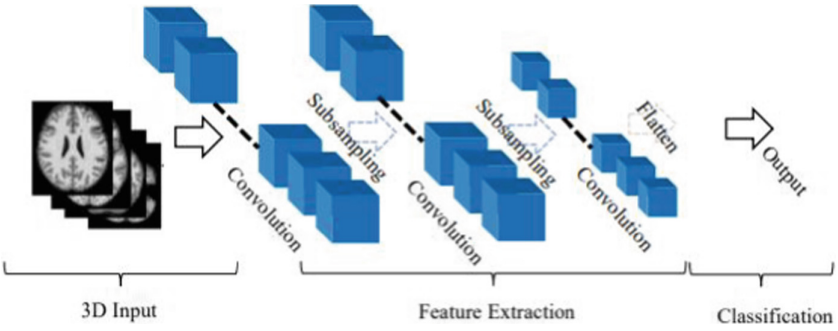

Fig. 2. 3D subject-level approach to dealing with volumetric input data [25].

4 Experiments and Results

This section presents the experiments for automatic diagnosis of Alzheimer's Disease (AD) using the ADNI database. The two CNN-based classifiers, introduced in the previous section, attempt to discriminate between Cognitively Normal (CN) and Alzheimer's Disease (AD) classes. The models were trained on a remote server supported by NVIDIA GeForce RTX 2080 Ti graphics cards using the Keras and the Tensorflow environments.

4.1 Dataset

The ADNI dataset consists of FDG-PET scans saved in the NII file format (or NIfTI), typically used for neuroimaging data. NIfTI stands for Neuroimaging Informatics Technology Initiative. The scans have been collected from different machines, with different resolutions, ranging from $128 \times 128 \times 35$ up to $400 \times 400 \times 144$ voxels, with the average resolution around $150 \times 150 \times 70$ voxels. The pixel intensity was normalized into the 0–255 interval and the images were cropped to a certain consistent resolution.

Data consists of 1355 total samples from witch 866 CN samples (63.91%), 489 AD samples (36.09%), 796 male patients (58.75%) and 559 female ones (41.25%). Data was divided into 1250 training samples and 105 testing samples (66 from the CN class and 39 from the AD class). The 1250 training samples were splitted into 10 folds for K-fold Cross Validation, 80 CN samples and 45 AD samples for each fold.

4.2 2D Slice-Level CNN for AD Diagnosis with PET Data

Google's Inception V3 was chosen as the 2D Slice -level CNN architecture pre-trained on ImageNet dataset (1000 classes, around 1.3 million data samples). The inception blocks are also known as the "mixed" blocks. Four variations of Inception V3 were trained - Mixed7, Mixed8, Mixed9, Mixed10, where the index means the number of the inception blocks. Only the last (fully connected) layer was fine tuned with the ADNI dataset. The 2D slices (Fig. 1) were obtained with the aid of OpenCV environment. We took care to group only slices belonging to the same subject and corresponding, approximately, to the same brain sections. The collages have uniformed dimensions, initially set at 512×512 pixels.

The results with respect to the four Mixed architectures are summarized in Fig. 3. The models were trained with SGD (Stochastic Gradient Descent) and the trainable parameters (in the last fully connected layer) were initialized with the pre-trained ImageNet weights. Although all models suffer from overfitting, the Mixed8 model outperforms the other architectures and, therefore, it was selected for further tuning.

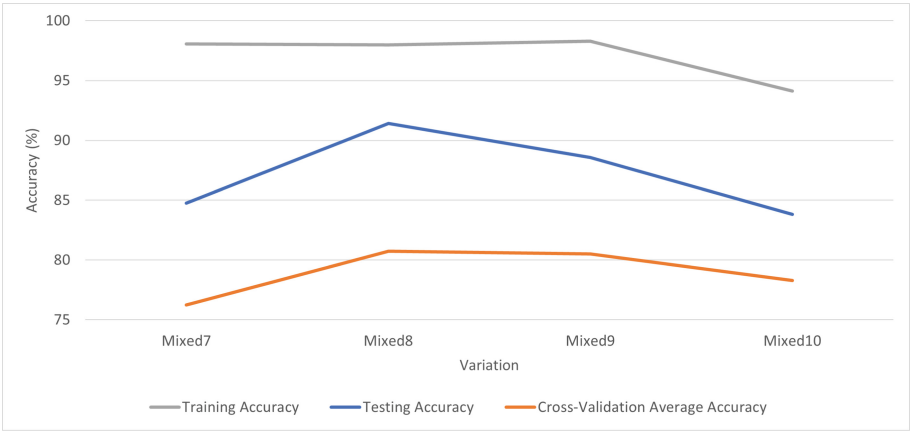

Fig. 3. 2D slice-level models: impact of the Inception V3 architecture (SGD optimizer; initialization with pre-trained ImageNet weights).

The importance of the optimization method and the impact of the parameter initialization were validated for all architectures. Figure 4 and Fig. 5 depict the results only for Mixed8 model. SGD was the most favourable optimizer and set up for the next experiments. Starting from the optimal parameters obtained at the pre-trained stage with the ImageNet reveals to be advantageous compared to random weights initialization.

Though the Mixed8 model reached a promising testing accuracy of 91.43% it still suffers overfitting. This problem was tackled by adding a dropout layer. Figure 6 depicts the classifier performance for a range of dropout rates. Note that 50% dropout rate appears to be a reasonable compromise between the overfitting and the fast convergence of the loss function as shown in Fig. 7. Smaller the dropout rate, faster the convergence, however more prone to overfitting.

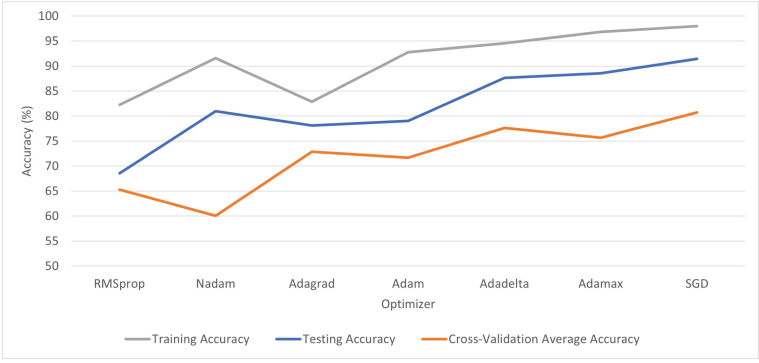

Fig. 4. 2D slice-level model: impact of the optimizer (Inception V3 mixed8; initialization with pre-trained ImageNet weights).

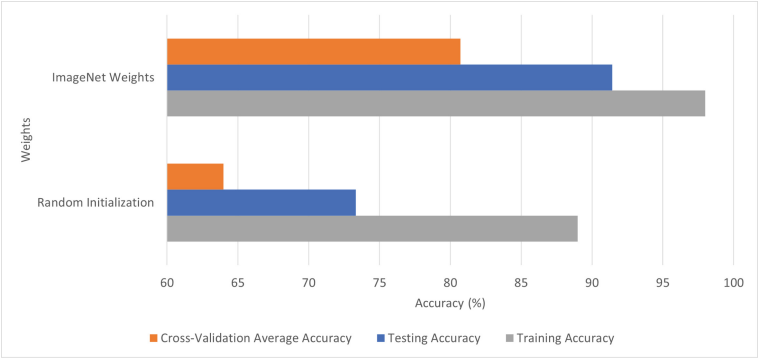

Fig. 5. 2D slice-level model: impact of the weight initialization (Inception V3 mixed8; SGD optimizer).

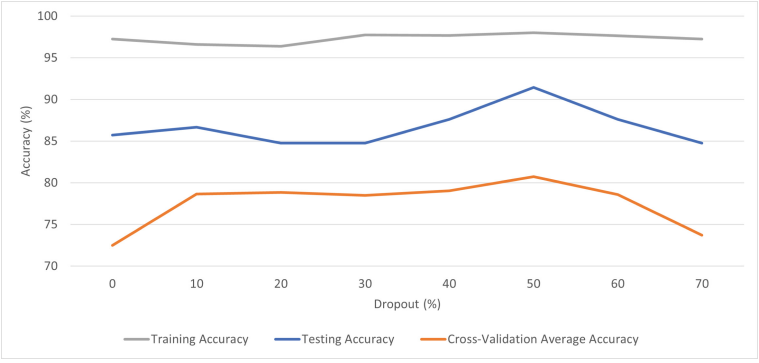

Fig. 6. 2D slice-level model: impact of dropout rates (Inception V3 mixed8; initialization with pre-trained ImageNet weights; SGD optimizer).

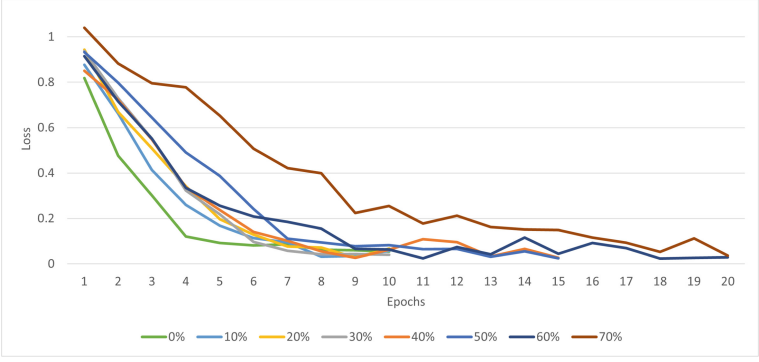

Fig. 7. 2D slice-level model: loss function trajectory for varying dropout rates (Inception V3 mixed8; initialization with pre-trained ImageNet weights; SGD optimizer).

4.3 3D Subject-level CNN for AD Diagnosis with PET Data

The CNN model used for this experiment was end-to-end designed and optimized. Exhaustive search for the optimal architecture is computationally infeasible. Instead, we selected a similar topology to the one proposed in [25]. The base structure consists of two groups of two 3D Conv layers and a 3D max-pooling layer, followed by a batch normalization and a flatten layer. The implementation code for the base architecture is shown in Fig. 8. Four variations of the base architecture (see Table 3) were trained with the binary cross-entropy loss function, SGD optimizer and PET images with dimension of $75 \times 75 \times 30$ voxels. The variations, basically, consist in changing the number of the Conv filters.

```
print("################### NETWORK ARCHITECTURE ###################")

input_format = Input((75, 75, 30, 1))

conv1 = Conv3D(filters=16, kernel_size=(3, 3, 3), activation='relu')(input_format)
conv2 = Conv3D(filters=8, kernel_size=(3, 3, 3), activation='relu')(conv1)

max_pool1 = MaxPool3D(pool_size=(2, 2, 2))(conv2)

conv3 = Conv3D(filters=16, kernel_size=(3, 3, 3), activation='relu')(max_pool1)
conv4 = Conv3D(filters=8, kernel_size=(3, 3, 3), activation='relu')(conv3)

max_pool2 = MaxPool3D(pool_size=(2, 2, 2))(conv4)

norm1 = BatchNormalization()(max_pool2)
flat1 = Flatten()(norm1)

fc1 = Dense(units=512, activation='relu')(flat1)

drop1 = Dropout(0.4)(fc1)

fc2 = Dense(units=128, activation='relu')(drop1)

drop2 = Dropout(0.4)(fc2)

fc3_out = Dense(units=1, activation='sigmoid')(drop2)
```

Fig. 8. 3D binary classification - base custom 3D-CNN architecture.

Table 3. 3D CNN - custom variations.

Model	Description
Custom 1	conv1(16 filters); conv 2(8 filters); conv3 (16 filters); conv4(8 filters)
Custom 2	conv1(8 filters); conv 2(16 filters); conv3 (8 filters); conv4(16 filters)
Custom 3	conv1(16 filters); conv 2(16 filters); conv3 (8 filters); conv4(8 filters)
Custom 4	conv1(8 filters); conv 2(8 filters); conv3 (16 filters); conv4(16 filters)

The results in terms of training, cross-validation and testing accuracy are depicted in Fig. 9. Custom4 model outperforms the other models and it is used in the next experiments. The impact of the 3D PET image resolution and the batch size were analysed as shown in Fig. 10 and Fig. 11. Based on these results, the Custom4 model trained with the $75 \times 75 \times 30$ voxels input image resolution and batch size 2 was considered as the optimal training configuration. Similarly to the 2D approach, the overfitting issue was tackled through the variation of the

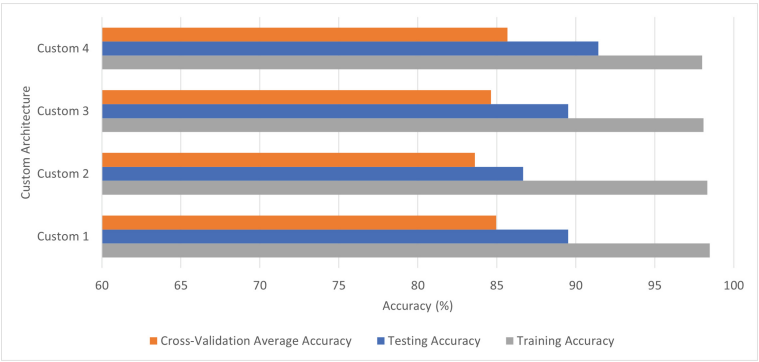

Fig. 9. 3D subject-level model: impact of the custom architecture ($75 \times 75 \times 30$ voxels; SGD optimizer; 0.001 learning rate; 50% dropout rate; batch size = 2).

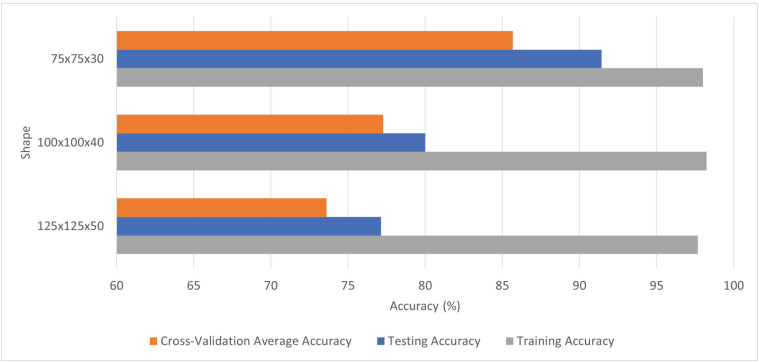

Fig. 10. 3D subject-level model - impact of the 3D PET image resolution (SGD optimizer; 0.001 learning rate; Custom4 model; 50% dropout rate; batch size = 2).

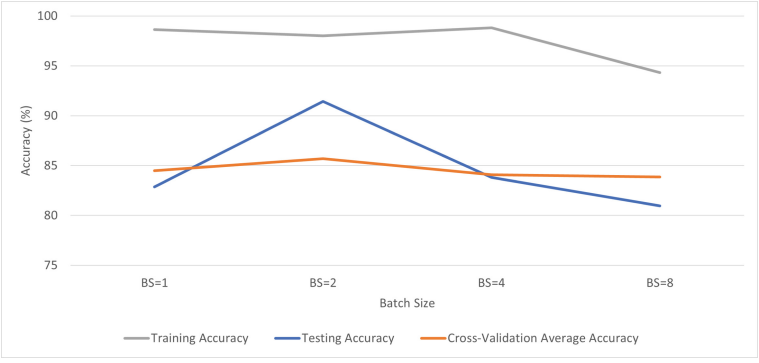

Fig. 11. 3D subject-level model - impact of the batch size ($75 \times 75 \times 30$ voxels; SGD optimizer; 0.001 learning rate; Custom4 model; 50% dropout rate).

dropout rate (see Fig. 12). The model struggles to converge for higher dropout rates, achieving a remarkable 91.43% testing accuracy for a 50% dropout rate (Fig. 13).

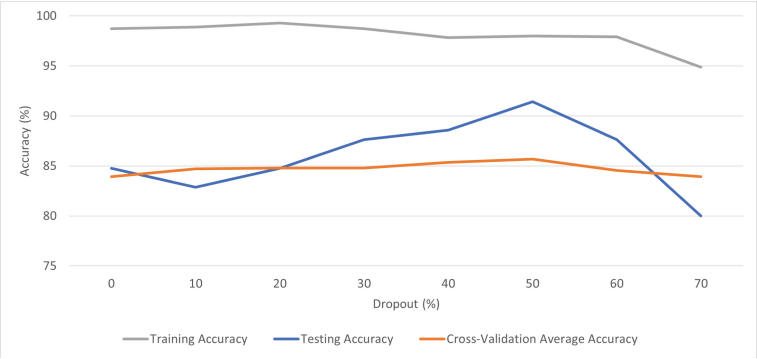

Fig. 12. 3D subject-level model - impact of the dropout rate ($75 \times 75 \times 30$ voxels; SGD optimizer; 0.001 learning rate; Custom4 model; batch size = 2).

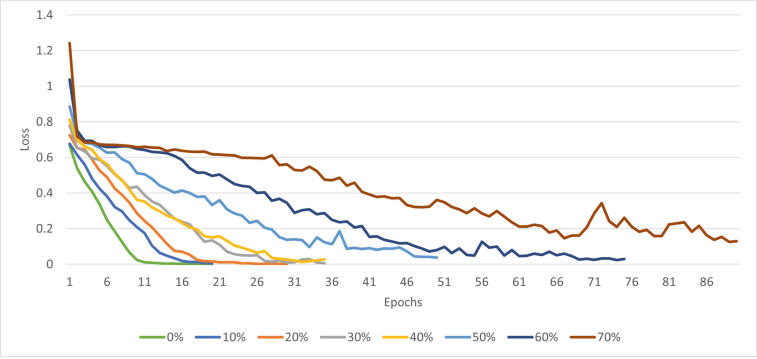

Fig. 13. 3D subject-level model: - loss function trajectory for varying dropout rate ($75 \times 75 \times 30$ voxels; SGD optimizer; 0.001 learning rate; Custom4 model; batch size = 2).

5 Conclusions

The primary objective of this paper was to study the potential of ^{18}F-FDG PET neuroimaging as a AD biomarker for classifying healthy versus AD patients. The first CNN model explores transfer learning with a pre-trained 2D Inception V3 model, as a typical solution in medical imaging. The second solution involves a custom 3D-CNN designed and trained from scratch. Both models achieved

competitive performance (testing accuracy above 91%), with scores above most of the referred works in Table 2. The custom 3D-CNN is computationally more attractive because it has less conv layers and, therefore, less number of parameters. Further to that, this study demonstrates that the 3D-CNN, provided with the FDG-PET data, is a promising brain imaging tools for AD diagnostics.

References

1. 2020 Alzheimer's disease facts and figures: Alzheimer's and Dementia (2020). https://doi.org/10.1002/alz.12068
2. Aderghal, K., Benois-Pineau, J., Afdel, K., Gwenaëlle, C.: FuseMe: classification of sMRI images by fusion of deep CNNs in 2D+e projections. In: ACM International Conference Proceeding Series (2017). https://doi.org/10.1145/3095713.3095749
3. Backstrom, K., Nazari, M., Gu, I.Y.H., Jakola, A.S.: An efficient 3D deep convolutional network for Alzheimer's disease diagnosis using MR images. In: Proceedings of the International Symposium on Biomedical Imaging (2018). https://doi.org/10.1109/ISBI.2018.8363543
4. Bozhkov, L., Georgieva, P.: Overview of deep learning architectures for EEG-based brain imaging. In: 2018 International Joint Conference on Neural Networks (IJCNN). IEEE (2018)
5. Bozhkov, L., Georgieva, P.: Deep learning models for brain machine interfaces. Ann. Math. Artif. Intell., 1175–1190 (2019). https://doi.org/10.1007/s10472-019-09668-0
6. Choi, H., Ha, S., Im, H.J., Paek, S.H., Lee, D.S.: Refining diagnosis of Parkinson's disease with deep learning-based interpretation of dopamine transporter imaging. NeuroImage Clinical (2017). https://doi.org/10.1016/j.nicl.2017.09.010
7. Choi, H., Jin, K.H.: Predicting cognitive decline with deep learning of brain metabolism and amyloid imaging. Behav. Brain Res. (2018). https://doi.org/10.1016/j.bbr.2018.02.017
8. Cole, J.H., et al.: Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. Neuroimage (2017). https://doi.org/10.1016/j.neuroimage.2017.07.059
9. Ding, Y., et al.: A deep learning model to predict a diagnosis of Alzheimer disease by using 18 F-FDG PET of the brain. Radiology (2019). https://doi.org/10.1148/radiol.2018180958
10. Ebrahimighahnavieh, M.A., Luo, S., Chiong, R.: Deep learning to detect Alzheimer's disease from neuroimaging: a systematic literature review. Comput. Methods Programs Biomed. (2020). https://doi.org/10.1016/j.cmpb.2019.105242
11. Havaei, M., et al.: Brain tumor segmentation with Deep Neural Networks. Med. Image Anal. (2017). https://doi.org/10.1016/j.media.2016.05.004
12. Huang, Y., Xu, J., Zhou, Y., Tong, T., Zhuang, X.: Diagnosis of Alzheimer's disease via multi-modality 3D convolutional neural network. Front. Neurosci. (2019). https://doi.org/10.3389/fnins.2019.00509
13. Kamnitsas, K., et al.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Med. Image Anal. (2017). https://doi.org/10.1016/j.media.2016.10.004
14. Klöppel, S., et al.: Automatic classification of MR scans in Alzheimer's disease. Brain (2008). https://doi.org/10.1093/brain/awm319

15. Li, F., Liu, M.: Alzheimer's disease diagnosis based on multiple cluster dense convolutional networks. Comput. Med. Imaging Graph. (2018). https://doi.org/10.1016/j.compmedimag.2018.09.009
16. Lian, C., Liu, M., Zhang, J., Shen, D.: Hierarchical fully convolutional network for joint atrophy localization and Alzheimer's disease diagnosis using structural MRI. IEEE Trans. Pattern Anal. Mach. Intell. (2020). https://doi.org/10.1109/TPAMI.2018.2889096
17. Lin, W., et al.: Convolutional neural networks-based MRI image analysis for the Alzheimer's disease prediction from mild cognitive impairment. Front. Neurosci. (2018). https://doi.org/10.3389/fnins.2018.00777
18. Liu, M., Cheng, D., Wang, K., Wang, Y.: Multi-modality cascaded convolutional neural networks for Alzheimer's disease diagnosis. Neuroinformatics, 295–308 (2018). https://doi.org/10.1007/s12021-018-9370-4
19. Liu, M., Zhang, J., Adeli, E., Shen, D.: Landmark-based deep multi-instance learning for brain disease diagnosis. Med. Image Anal. (2018). https://doi.org/10.1016/j.media.2017.10.005
20. Marcus, C., Mena, E., Subramaniam, R.M.: Brain PET in the diagnosis of Alzheimer's disease (2014). https://doi.org/10.1097/RLU.0000000000000547
21. Punjabi, A., Martersteck, A., Wang, Y., Parrish, T.B., Katsaggelos, A.K.: Neuroimaging modality fusion in Alzheimer's classification using convolutional neural networks. PLoS ONE (2019). https://doi.org/10.1371/journal.pone.0225759
22. Qiu, S., Chang, G.H., Panagia, M., Gopal, D.M., Au, R., Kolachalama, V.B.: Fusion of deep learning models of MRI scans, Mini-Mental State Examination, and logical memory test enhances diagnosis of mild cognitive impairment. Alzheimer's Dement. Diagn. Assess. Dis. Monit. (2018). https://doi.org/10.1016/j.dadm.2018.08.013
23. Senanayake, U., Sowmya, A., Dawes, L.: Deep fusion pipeline for mild cognitive impairment diagnosis. In: Proceedings of the International Symposium on Biomedical Imaging (2018). https://doi.org/10.1109/ISBI.2018.8363832
24. Shmulev, Y., Belyaev, M.: Predicting conversion of mild cognitive impairments to Alzheimer's disease and exploring impact of neuroimaging. In: Stoyanov, D., et al. (eds.) GRAIL/Beyond MIC 2018. LNCS, vol. 11044, pp. 83–91. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00689-1_9
25. Singh, S.P., Wang, L., Gupta, S., Goli, H., Padmanabhan, P., Gulyás, B.: 3d deep learning on medical images: a review (2020). https://doi.org/10.3390/s20185097
26. Stier, N., Vincent, N., Liebeskind, D., Scalzo, F.: Deep learning of tissue fate features in acute ischemic stroke. In: Proceedings - 2015 IEEE International Conference on Bioinformatics and Biomedicine, BIBM 2015 (2015). https://doi.org/10.1109/BIBM.2015.7359869
27. Valliani, A., Soni, A.: Deep residual nets for improved Alzheimer's diagnosis. In: ACM-BCB 2017 - Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics (2017). https://doi.org/10.1145/3107411.3108224