

Preclinical Stage Alzheimer’s Disease Detection Using Magnetic Resonance Image Scans

Fatih Altay^{1*}, Guillermo Ramón Sánchez^{1*}, Yanli James², Stephen V. Faraone², Senem Velipasalar¹, Asif Salekin^{1†}

¹ Syracuse University

² SUNY Upstate Medical University

faltay@syr.edu, grsanche@syr.edu, zhangy@upstate.edu, sfaraone@childpsychresearch.org, svelipas@syr.edu, asalekin@syr.edu

Abstract

Alzheimer’s disease is one of the diseases that mostly affects older people without being a part of aging. The most common symptoms include problems with communicating and abstract thinking, as well as disorientation. It is important to detect Alzheimer’s disease in early stages so that cognitive functioning would be improved by medication and training. In this paper, we propose two attention model networks for detecting Alzheimer’s disease from MRI images to help early detection efforts at the preclinical stage. We also compare the performance of these two attention network models with a baseline model. Recently available OASIS-3 Longitudinal Neuroimaging, Clinical, and Cognitive Dataset is used to train, evaluate and compare our models. The novelty of this research resides in the fact that we aim to detect Alzheimer’s disease when all the parameters, physical assessments, and clinical data state that the patient is healthy and showing no symptoms.

Introduction

Alzheimer’s disease (AD) is a type of brain disease, which is degenerative, and its symptoms worsen over the years. According to the Alzheimer’s Association¹, 5.8 million Americans are living with AD. Statistics have also shown that annual AD related death rates by age are increasing every year (vital statistics system 2019). According to (Reiman et al. 2012), AD could begin twenty years or more before symptoms are perceptible, and between one and six years with changes in the brain that are unnoticeable to the person affected. As the disease evolves, more and more brain neurons stop functioning, lose connection, even die. At first, AD affects the entorhinal cortex and hippocampus brain regions that are involved in memory (Zott et al. 2018). Eventually, it affects the cerebral cortex, which is responsible for language, reasoning, and social behavior, and many other areas of the brain.

This pathology, according to the Alzheimer’s Association¹, is the most common cause of dementia, especially among

older people. AD accounts for 60-80% of the total dementia cases (Manzano et al. 2009). AD significantly reduces the quality of life (Logsdon et al. 1999) and life expectancy (Rice et al. 2001) of patients, and is considered as the most expensive disease in USA (Hurd et al. 2013). When the economic impact of this disease is considered, just in 2019, the total combined payments from all AD patients are estimated at \$290 billion¹.

While there is no cure for Alzheimer’s disease or a way to stop its progression after diagnosis, and the treatment of AD is still an open research question, there are drug and non-drug options that may help treat symptoms. Studies have demonstrated that early stage intervention of AD can significantly impact the degeneration process, and treatment of symptoms (Kang et al. 2019; Golde, DeKosky, and Galasko 2018; Crous-Bou et al. 2017). The early detection of AD through conventional MRI scanning will facilitate effective and in-time interventions/treatments, that would expand the life expectancy and quality of life of patients.

National Institute on Aging has defined three stages of AD: (1) Preclinical, when the patients do not exhibit any symptoms, but brain neuronal structure has started to deteriorate; (2) Mild cognitive impairment (MCI), when patients start to exhibit cognitive impairments, but still can perform all activities of daily living (ADL); and (3) Alzheimer’s dementia, when symptoms of dementia are severe enough to interfere with ADLs.

With the advancement of MRI technology (Lohrke et al. 2016), and the recent development of deep learning based computer vision approaches, several studies have addressed the detection of AD in MCI and dementia stages (stage 2 and 3) from MRI brain scans (Cheng et al. 2017; Böhle et al. 2019). While these studies have improved our understanding of Alzheimer’s disease, they do not contribute in early stage interventions, since stage 2 and 3 are accurately identifiable through clinical diagnosis. Accurate detection or indication of preclinical AD is a major interest in the medical community (Schindler et al. 2017; Pernecky 2018). However, to the best of our knowledge, existing studies have not ad-

*Equal Contribution

†Corresponding author

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.alz.org/media/documents/alzheimers-facts-and-figures-2019-r.pdf>

addressed the challenge of preclinical AD detection from MRI brain scans yet.

Recent reports on AD suggest that (Sperling, Jack, and Aisen 2011) change in brain may be evident 20 years before the stage of dementia (stage 3), and that there is already substantial neuronal loss by the stage of mild cognitive impairment (MCI) (stage 2). Hence, the goal of this study is the development of an effective machine learning approach that can identify the latent patterns due to preclinical AD from MRI brain scans, which can significantly improve intervention and treatment of AD patients.

MRI brain scans are 3 dimensional (3D) image representation of the brain structure. Several sequential classifiers such as, 3D Convolutional Neural Network (CNN) and 3D recurrent visual attention (RVN) model (Wood, Cole, and Booth 2019), have been implemented to detect disease related patterns from the 3D brain scans. Recently, Transformer (Girdhar et al. 2019) model has been demonstrated to outperform all the existing sequential classifiers.

In this study, we employ and adapt two attention network models (3D recurrent visual attention model (Wood, Cole, and Booth 2019) and Transformer model (Girdhar et al. 2019)) to our problem of preclinical AD detection from 3D MRI brain scan data, and compare their performances with a baseline model that is based on 3D CNNs. We evaluated these approaches for differentiating individuals with ‘preclinical AD’ from ‘others’. The ‘others’ class includes healthy individuals or patients suffering from other dementia problems not related with AD. According to our evaluation, the 3D CNN model achieves an F1 score of 0.83 and accuracy of 88.34%, the 3D RVN model achieves an F1 score of 0.87, and 90.65% accuracy, and the developed transformer model achieves an F1 score of 0.90 and 91.18% accuracy in preclinical (i.e., prior to MCI and dementia stage) Alzheimer disease detection. Experimental results demonstrate that, by using the MRI images effectively, it is possible to detect preclinical stage Alzheimer’s disease with a very promising accuracy.

Related Work

Deep learning algorithms perform very well in identifying complex structures in high dimensional data, and this is why there is a very rich literature on detecting diseases using magnetic resonance images (MRI) and positron emission tomography (PET). (Korolev et al. 2017) studied the presence of Alzheimer’s disease using two different 3D CNN approaches, namely VoxCNN and residual neural networks (RNN), on MRI data. Also aiming for binary classification, Alzheimer’s/Non Alzheimer’s, (Cheng et al. 2017) used two 3D CNNs on MRI and PET scans, combining them with a fully connected layer and a softmax classification.

Over the past few years, there has been a trend in applying attention based approaches to medical problems. With these kind of models, we can not only detect a certain pathology, but also represent which part of the data is more important to make accurate predictions. (Schlemper et al. 2019) proposed a novel method, referred to as attention gated networks, for medical image analysis. This algorithm learns to focus on target structures, and can be used for leveraging

certain regions for classification purposes. With a similar goal, (Wood, Cole, and Booth 2019) proposed Neuro-Dram, a 3D recurrent visual attention, for explainable neuroimaging classification. (Böhle et al. 2019) used a layer-wise relevance propagation (LRP) to visualize CNN decisions based on MRI data. This algorithm attributes relevance to every input node, and studies the contribution of each node. It can output a heat map to highlight the most informative parts of every input image.

Early detection of AD can be critical when developing an optimal treatment for each patient (Yu et al. 2009). (Kang et al. 2019) showed how early stage AD patients, improved their general cognitive abilities after 12 weeks of paper-based cognitive training. Along these lines, there has been a lot of work trying to differentiate between mild cognitive impairment (MCI) and Alzheimer’s disease (AD). According to (Markesbery 2010), MCI causes a slight and measurable decline in cognitive abilities and it is the earliest clinically detectable stage before AD. Suffering from MCI increases the risk of developing AD. (Gupta et al. 2019) proposed a four-class SVM classifier: AD, MCI stable (patients with MCI who do not develop AD), MCI converted (patients whose MCI develop into AD), and healthy patients. Similarly, to study MCI to AD conversion, (Lee et al. 2019) proposed a multi-modal recurrent neural network. In this case, they used MRI, demographic information, cerebrospinal fluid (CSF) biomarkers and cognitive performance reports as their inputs to the GRU units. With a slightly different approach, (Islam and Zhang 2017) used a model inspired by Inception-V4 network (Szegedy et al. 2017) to detect between non demented, very mild dementia, mild dementia and moderate dementia.

The goal of all the discussed approaches is detecting AD (stage 3) or predicting MCI cases (stage 2) that will develop AD, i.e. existing studies focus on stages 2 and 3. To the best of our knowledge, there is no previous work on detecting preclinical stage AD (stage 1) when all the indicators, including the clinical assessments by doctors, ensure that the patient is healthy and no symptoms of the disease are present. Our main goal is to detect future AD from the latent brain scan patterns even before MCI develops, when the disease is in a preclinical stage.

Dataset

In this work, we employ the recently published longitudinal neuroimaging, clinical and cognitive dataset, called OASIS-3 (LaMontagne et al. 2019). It consists of MRI and PET imaging from 1098 individuals collected across several studies over the course of 15 years. There are 605 cognitively normal adults and 493 individuals at different stages of cognitive decline. Ages of the participants range from 42 to 95 years. The dataset contains over 2000 MRI sessions.

For every patient, the aging and disability resource center (ADRC) clinical data is available. This data includes patient’s age at entry, height, weight, and clinical dementia rating (CDR). Diagnoses for this datatype include “cognitively normal”, “AD dementia”, “vascular dementia” and factors that could be contributing such as vitamin deficiency, alcoholism, and mood disorders. The goal of this study is to iden-

tify potential AD patients even when they were clinically diagnosed as “cognitively normal” from the latent patterns of their brain scans. Hence, in this work, we are only using brain imaging and the affiliated clinical diagnoses, matching each scan with the closest clinical diagnosis available.

The presented classifiers perform a binary classification task: differentiating preclinical AD individuals from others. More specifically, our presented binary classifiers use ‘0’ for the patients diagnosed as healthy or with other non-Alzheimer’s related pathologies, and ‘1’ for the patients with ‘preclinical stage AD’. We define a ‘preclinical stage AD’ patient as a person who is currently diagnosed by a doctor as healthy but we know that in the future he or she will develop AD. For this, we label each MRI session independently, matching every clinical data that we have with the closest MRI session available.

In this dataset, MRI brain scans are available from two different views, the axial or horizontal plane, and sagittal or longitudinal plane (Park et al. 2010). We decided to use the axial plane since only a smaller portion of the individual’s data contain sagittal scans. Since not all the images had the same resolution, as part of our image pre-processing, we re-sized and normalized all the images.

One of the initial challenge we encountered was the significant class imbalance. More specifically, for class ‘0’ (healthy subjects), there were 2181 scans, whereas for class ‘1’ (patients with preclinical stage AD), we had only 176 scans. To address this issue, we down-sampled our class ‘0’ and over-sampled the class ‘1’, by randomly rotating and mirroring some of the images (as proposed in (Ando and Huang 2017) and (Zhou 2020)). We also implemented a balanced sampling process. This method creates a sampler in the data loader based on the number of images of each class that yields the next index/key set to fetch. This is useful to ensure that we have an even number of labels on each training batch.

Proposed Approach

As mentioned above, existing works addressed stage 2 and 3 AD detection. One of the significant differences of our work, compared to the state-of-the-art, is that we address the more challenging problem of preclinical AD (stage 1) detection. To ensure the latent brain pattern extraction for preclinical AD detection task, we only considered the 3D brain scans as input. This makes the classification task more challenging, since we cannot adapt the models to different factors that could be critical when predicting AD, such as the age or the sex of the patient. To address this challenging problem, we adapt and employ two different attention mechanisms, namely 3D recurrent visual attention model, and attention transformer. We compare our two attention mechanisms with a baseline model, which is based on 3D CNNs.

Baseline - 3D CNN Model

The baseline model is based on a 3D CNN model, which was initially used for video classification tasks². This model uses 3D kernels and channels to convolve video input, where

²<https://github.com/HHTseng/video-classification>

the videos are viewed as 3D data (2D images over time dimension). For our baseline model, we stack all the images in a brain scan, turn them into 3D input data, and then feed it to the network. The model we developed consists of five convolutional layers and three fully connected layers. Each convolution layer is followed by batch normalization, ReLU, dropout, and pooling layers. We consider this model as our baseline.

3D Recurrent Visual Attention Model

For our first model shown in Fig. 1, we employ a recently proposed 3D recurrent visual attention model, which is tailored for neuroimaging classification (Wood, Cole, and Booth 2019) and focuses on already developed AD detection task. This model uses a recurrent attention mechanism that tries to find relevant locations of brain scan indicative of AD. The model consists of an agent that is trained with reinforcement learning. It is built around a two-layer recurrent neural network (RNN). At each timestamp, the agent receives a small portion of the entire image, which is a glimpse, centered around a position l , and decides which location to select at the next timestamp. After a fixed number of steps, a classification decision is made. The aim of using an agent is to maximize the rewards along the timestamps, and then decide to attend the most informative regions of the images. We define our reward, r_t , as 1 for all timestamp t if the classification is correct, or 0 if it is not. Overall, the model consists of four different networks: the glimpse network, the recurrent network, the location network and the classification network.

The glimpse network takes a small 3D image fraction x_t , and its location coordinate l_t as input, and outputs a vector g_t . It generates a representation of the glimpse (i.e., 3D image fraction) summarizing the ‘what’ g_{x_t} , and ‘where’ g_{l_t} information. The glimpse network consists of 3D convolutional layers (with batch normalization and max pooling) that generates the ‘what’ representation, and a single-layer fully connected layer that converts the location coordinated to the ‘where’ representation. Final g_t is obtained by an element-wise multiplication of these representations:

$$g_t = g_{x_t} \odot g_{l_t} \quad (1)$$

The recurrent network is used to obtain the agent’s internal representation encapsulating the information extracted from past timestamps. This network consists of two stacked LSTM units (Hochreiter and Schmidhuber 1997). At each timestamps, hidden-layer representation generated by the LSTM units is fed into the location network to obtain the next timestamp’s glimpse location, l_{t+1} . At the last timestamp (i.e., last iteration of the LSTM sequential analysis) the hidden layer representation is fed into the classification network.

The location network consists of a single-layer fully connected layer, which maps $r_{(t)}^2$ to a 3D vector in the range $[-1, 1]$ which is an isotropic 3D normal distribution. Next location l_{t+1} is then produced by sampling from this distribution.

Finally, **the classification network** consists of a single fully connected layer with a sigmoid activation function,

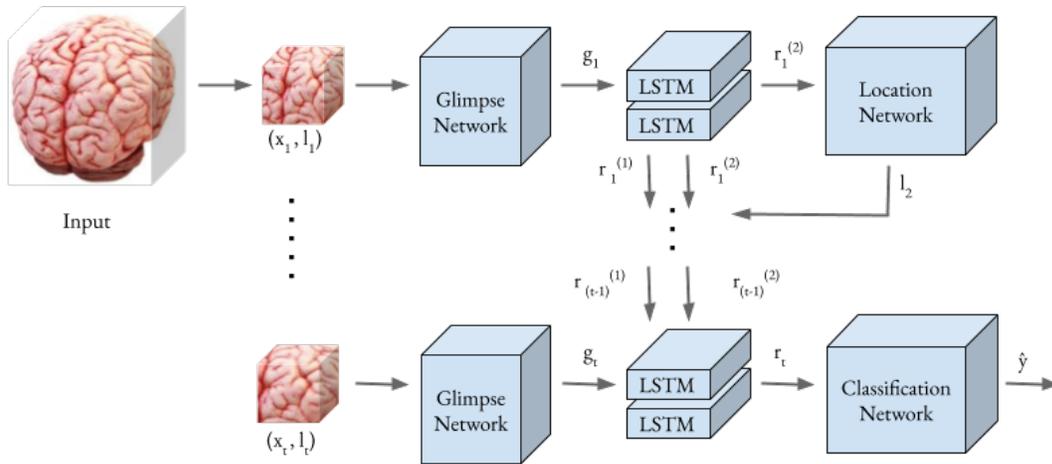


Figure 1: Overview of 3D Recurrent Visual Attention Model

which is used for binary classification. As previously stated, the input to this network is the final timestamp hidden layer representation of the LSTM unit.

Total number of timestamps/iteration T in the recurrent network, and the glimpse size are hyper-parameters. In our evaluation we identified $T = 6$, and the glimpse size (3D image fraction) $40 \times 40 \times 40$ as the beneficial value, based on the grid search on possible values.

Attention Transformer

As our second model, we employed a transformer network for the task of preclinical AD detection. Transformer models have been used for different tasks such as human action recognition from videos (Girdhar et al. 2019) and text translation (Vaswani et al. 2017). Although transformer networks have been used for other tasks and applications, we firmly believe that this is the first work that employs a transformer network on MRI images of brain for preclinical stage Alzheimer’s disease detection. Slices from a brain scan are fed to the network, and the network is expected to detect whether any sign of dementia is observable or not, even the subject is showing no signs nor symptoms of the disease yet.

Implemented Transformer model consists of a base network and a head network similar to (Girdhar et al. 2019). Base network extracts feature-representations from each of the brain MRI image slices, and the head network generates the binary inferences (i.e., preclinical stage AD or others). Both the base and the head networks are described along with the modifications we have made in the following subsections.

Base Network. Since we deal with binary classification for preclinical stage AD detection, we use a VGG16 network (Simonyan and Zisserman 2015) as our base network. We extract frames from brain MRI scans (96 frames) and feed them into the base network. Due to the nature of the base network, it accepts $224 \times 224 \times 3$ images, where 224×224 represents image height and weight, and 3 represents the channel size. On the other hand, slices from brain MRI scans differ

in size, and are gray scale, which has only one channel. After resizing our input images to 224×224 , the first modification we do for the base network is adding an extra convolutional layer at the beginning to get the desired input for the base network. We set `in_channels=1`, `out_channels=3`, `kernel_size=3`, `stride=1`, `padding=1` and `dilation=1` for this new convolutional layer.

We have experimented with two different approaches, namely training the base network from scratch and using a pre-trained model (trained on conventional RGB images). As detailed in the experiments section, we obtained better results when the base network was trained from scratch. Instead of stacking images, (compared to 3D CNN and 3D recurrent visual attention model), we feed each of the brain scan images/frames in a sequence, like a video. The base network portion of the Fig. 2 illustrates the network described so far. We then send the output of the base network to the head network.

Transformer Unit. The original transformer architecture was proposed in (Vaswani et al. 2017) for sequence to sequence tasks to overperform recurrent models. It is done by selecting a feature frame and comparing it with all features in a sequence in order to compute attention. To do that, features are mapped to a query (Q) and memory (K for key and V for value) embeddings using linear projections. Since the original transformer is designed for language to language translation, Q is the word that is being translated and K and V are the linear projections of the generated input and output sequences.

On the other hand, our transformer network consists of a positional encoder, three head units, shown as Block Head Units in Fig. 2, and a classifier. This unit takes the brain MRI image sequence features and also the positional embeddings for attention proposal regions and maps them into query (Q) and memory features (K, V). Query (Q) represents the region with signs of preclinical stage AD. Frames around that region are projected into K and V. Block Head Units as in Fig. 2 process the query and memory embeddings to update

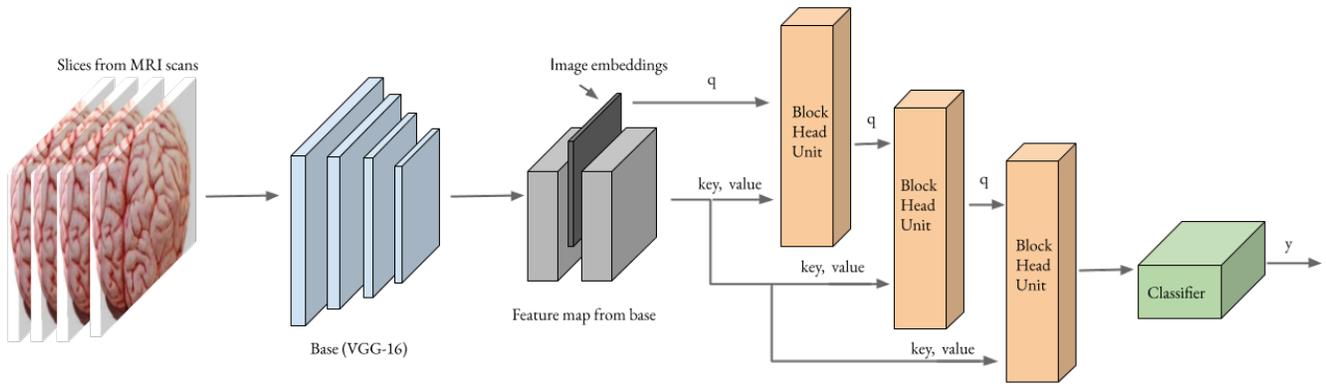


Figure 2: Overview of Attention Transformer Model

the query vector. By doing so, they aggregate the information over the brain MRI scans to classify whether the given sequence belongs to preclinical stage AD class or not.

We evaluated using different optimizers, and SGD (Robbins 2007) was the most beneficial choice. We set learning rate as $1e^{-4}$, momentum as $9e^{-4}$ and nesterov as True (Sutskever et al. 2013). We use cross-entropy loss (Korolev et al. 2017) for our classifier.

Experiments

We first split our dataset into three person-disjoint subsets: training, validation and test subsets. When trying to detect a certain pathology using patients' scans, it is crucial to split the data correctly. Different scans from the same person cannot appear among different datasets. If so, there is a risk that the model gets trained to correctly label those repeated patients, and it will perform poorly on scans from never-before-seen patients. This is why we decided to use person disjoint datasets. Each participant's MRI scan consists of 256 images. After the downsampling and oversampling process previously described, we used 65% of the available data for training, 20% for validation and 15% for testing.

It should be emphasized that the goal of this study is to detect scans, which were confirmed to be healthy by the doctors at the time of the scan, and which we know that in the future will develop Alzheimer's disease. Hence, we developed binary classifiers with preclinical stage Alzheimer's patients labeled as class '1', and the healthy patients or patients suffering from other dementia problems not related with Alzheimer's labeled as '0'. The evaluations discussed in this section follow the data labeling scheme shown in Fig. 3. In line with the study goal, we exclude the scans of patients, who have already developed and been diagnosed with AD, from the 'other' category (i.e., class 0), in order to focus on better differentiating between healthy and pre-clinical AD patients. As previously stated, to the best of our knowledge, there is no previous work on detecting pre-clinical stage AD.

For each model, we calculate the confusion matrix, F1 score, false negative rate, precision, recall, and accuracy. One of the best metrics to compare models when dealing

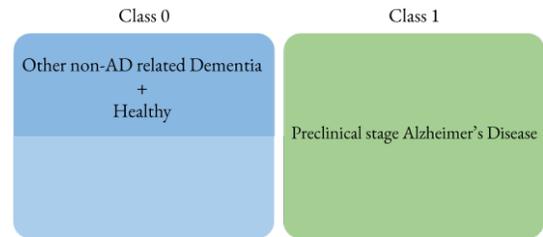


Figure 3: Overview of our data labeling approach

with an imbalanced classification problem is the F1 score³. In the following subsections, we will present the obtained results for each model, and compare them in terms of the F1 score. Also, since our main task is being able to detect possible AD patients when they are still healthy, the false negative rate will also be considered as an important comparative metric.

Results for the 3D CNN model. We used our 3D CNN model as our baseline. The best version of this model is trained with a learning rate of $1e^{-5}$ for 50 epochs. To avoid overfitting the training data, we chose AdamW as our optimizer with a weight decay of 0.1. We used cross-entropy (Korolev et al. 2017) as our loss function. With these configurations, we obtained the confusion matrix shown in Table 1.

		Actual	
		Healthy / other dementia	Preclinical AD
Predicted	Healthy / other dementia	99	15
	Preclinical AD	4	45

Table 1: Confusion matrix for the baseline model.

According to the evaluation, the 3D CNN model achieves

³<https://towardsdatascience.com/metrics-for-imbalanced-classification-41c71549bbb5>

an F1 score of 0.83 on preclinical stage AD detection. As shown in Table 1, we have a false negative (FN) rate of 25%. By taking a closer look at some of the patients wrongly labeled as healthy or other dementia when they were pre-clinical AD patients, we identified that they are the patients whose clinical data is not very consistent. More specifically, there are some patients who are diagnosed with AD at some point of the data collection process, and then they recover their cognitive abilities. Therefore, these patients are diagnosed as “Cognitive normal” for the rest of their life. We find that this type of patients are the hardest to correctly classify. Additionally, OASIS 3 dataset contains brain scans from patients with non-Alzheimer (non-AD) related dementia, that 3D CNN baseline model finds difficult to classify correctly. On the other hand, the baseline model performs well in classifying subjects who, once diagnosed with AD dementia, do not fully recover their cognitive abilities.

As an example, the clinical data of patient ID 30205 (on OASIS 3 dataset) is shown in Table 2. The patient was clinically diagnosed as “Cognitively normal” on day 0000, 0406, and 0773. OASIS 3 Dataset contains a complete MRI brain scan of the patient performed at day 61. Our baseline 3D CNN model can detect ‘preclinical stage AD’ from these scans performed on day 61. Thus, the baseline algorithm is able to detect ‘preclinical stage AD’ 1064 days (1125 - 61 = 1064) before the patient is diagnosed with uncertain dementia. If we just consider when the patient is actually diagnosed with Alzheimer’s disease, we are able to detect that pathology 1,776 days (1837-61) before it is diagnosed.

Day	Diagnose
0000	Cognitively normal
0406	Cognitively normal
0773	Cognitively normal
1125	uncertain dementia
1460	uncertain dementia
1837	Alzheimer’s disease dementia

Table 2: Clinical data for patient 30205.

Results for the 3D recurrent visual attention model.

With the attention mechanism models incorporated, our goal is to outperform the baseline. For the 3D recurrent visual attention (3D RVN) model, the best performance is achieved with a learning rate of $1e^{-4}$, trained for 200 epochs using AdamW. We also evaluated different glimpse sizes (i.e., image fraction size). More specifically, we experimented with $20 \times 20 \times 20$, $40 \times 40 \times 40$, and $60 \times 60 \times 60$. The best results presented in Table 3 were obtained with the $40 \times 40 \times 40$ glimpse size. According to the evaluation, the 3D RVN model achieves an F1 score of 0.87 for ‘preclinical stage AD’ detection, which is 5% higher compared to the baseline 3D CNN model. Additionally, according to the results, the false negative (FN) rate is 19.2%, which is lower than the 3D CNN model.

As an example, the clinical data of patient ID 30025 (on OASIS 3 dataset) is shown in Table 4. The patient was clinically diagnosed as “Cognitively normal” from day 0000 to

		Actual	
		Healthy / other dementia	Preclinical AD
Predicted	Healthy / other dementia	84	10
	Preclinical AD	3	42

Table 3: Confusion matrix for the 3D RVN model.

day 2608, and diagnosed as “AD dementia” on day 2933. OASIS 3 Dataset contains two complete MRI brain scans of the patient performed on day 0210 and 2298. Our 3D RVN model correctly infers ‘preclinical stage AD’ by taking the brain scan taken on day 0210 as input. This result demonstrates that the presented 3D RVN can detect AD 2723 days before it is clinically diagnosed by doctors.

Day	Diagnose
0000	Cognitively normal
0359	Cognitively normal
0751	Cognitively normal
1106	Cognitively normal
1547	Cognitively normal
1915	Cognitively normal
2247	Cognitively normal
2608	Cognitively normal
2933	Alzheimer’s disease dementia

Table 4: Clinical data for patient 30025.

One of the important benefits of this model is that not only we can detect AD before any symptoms, but also we can plot the part of the data that the algorithm interprets as more important when making accurate decisions. In our implementation, the initial (timestamp 0) glimpse location l_t is at the center of the brain. After that, the glimpse moves in a direction that maximizes the reward function.

Figure 4 shows the extracted glimpse locations by the 3D RVN model for a random participant with preclinical AD. The figure demonstrates how the 3D RVN agent is exploring the brain regions. The green point shows the first glimpse, which is located at the center of the brain at the first timestamp. Since we are using a total of six timestamps ($T = 6$) in the current 3D RVN, there are six different areas where the model agent focuses on to make an inference. By plotting the glimpse’s location, we can visualize which parts of the brain are more informative when detecting preclinical stage Alzheimer’s disease.

By examining the testset brain scans, and the areas where the 3D RVN model is paying more attention (producing location l_t to extract glimpse/3D image fraction), we identified that it is focusing on areas like: locus coeruleus, hippocampus, entorhinal cortex, and amygdala, parts of the brain that are important for memory and very relevant when detecting MCI (Bauer, Cabral, and Killiany 2018). Another key brain region for detecting AD is the lateral ventricles (Ertekin et al. 2016), and the 3D RVN model is focusing on that area

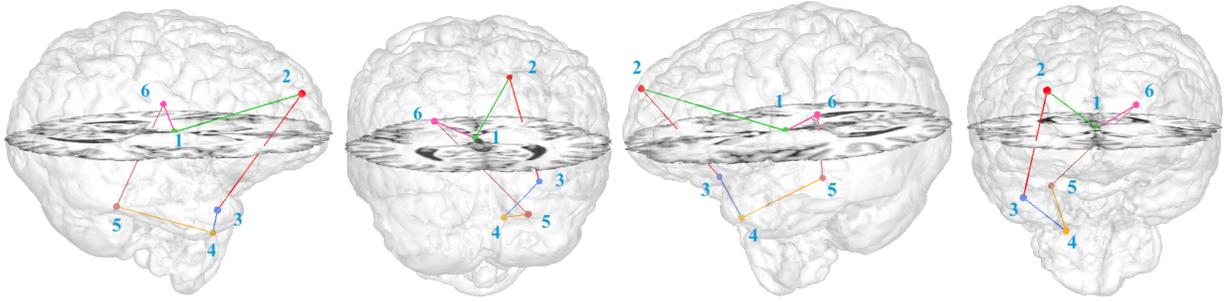


Figure 4: Trajectory taken by the glimpse algorithm. The green point is the first location, followed by the red, blue, yellow, brown, and pink points. We draw a line between them for visualization purposes.

as well.

Results for the attention transformer model. The transformer model is trained in a different way compared to other networks. First, instead of feeding stacked 3D images to the network, we feed images in a sequence. As mentioned in the Proposed Approach section, we use a VGG16 (Simonyan and Zisserman 2015) network as our base to extract features. We tried a pre-trained VGG16 model to see the effects of pre-trained weights, and also trained the network from scratch. As shown in Table 5, the transformer model when trained from scratch achieved 4.6% higher F1 score.

	Pre-trained Transformer	Our transformer
Accuracy	88.82%	91.18%
F1 Score	0.86	0.90
Precision	0.91	0.85
Recall	0.76	0.96

Table 5: Comparison between pre-trained transformer and our transformer.

Due to memory limitations, we are not able to use all frames from brain MRI scans. We started by selecting 48 frames from the middle and then went up to 96 frames, which is the highest amount we can select with our current GPU configuration. We achieved the best results with 96 frames, and trained the network with same amount of images from every scan. We also trained the network for different epochs. Due to the nature of the transformer networks, we need to select higher number of epochs for an optimal training. Thus, we trained our transformer model for 200 epochs.

As for the optimizer, we started with Adam (Kingma and Ba 2015), and set the learning rate (lr) to $1e^{-4}$ and *amsgrad* parameter to true. Then, we trained the network with AdamW with $lr = 1e^{-4}$ and *amsgrad* set as true. Finally, we trained the transformer network with SGD (Robbins 2007), and set $lr = 1e^{-4}$, momentum as $9e^{-4}$, and nesterov as True (Sutskever et al. 2013). Among all the optimizers, we obtained better results, presented in Table 6, with SGD.

According to the evaluation results shown in Table 6, the transformer model achieves an F1 score of 0.90, not only

		Actual	
		Healthy / other dementia	Preclinical AD
Predicted	Healthy / other dementia	87	12
	Preclinical AD	3	68

Table 6: Confusion matrix for the transformer model.

improving the baseline model but also outperforming the 3D RVN model. In terms of the false negative (FN) rate, the transformer model also achieves the lowest rate with 15%.

The transformer model also achieves the best performance in terms of the earliest detection. The clinical data of patient ID 30557 is shown in Table 7. The patient was clinically diagnosed as ‘‘Cognitively normal’’ from day 0000 to day 3816, and first diagnosed as ‘‘AD dementia’’ on day 4222. OASIS 3 Dataset contains two complete MRI brain scans of the patient performed on days 1448 and 2185. Transformer model infers ‘preclinical stage AD’ by taking the brain scan taken on day 1448 as input. This means that the algorithm is detecting ‘preclinical AD’ 2774 days ($4222 - 1448 = 2774$) before the patient is clinically diagnosed.

Day	Diagnose
0000	Cognitively normal
0363	Cognitively normal
0749	Cognitively normal
1076	Cognitively normal
1464	Cognitively normal
1980	Cognitively normal
2378	Cognitively normal
3093	Cognitively normal
3452	Cognitively normal
3816	Cognitively normal
4222	Alzheimer’s disease dementia
4586	Alzheimer’s disease dementia

Table 7: Clinical data for patient 30025.

Comparison between models. In Table 8, we compare the three models by measuring accuracy, precision, recall,

Model	Accuracy	F1 Score	Precision	Recall
Baseline	88.34%	0.83	0.92	0.75
RVN	90.65%	0.87	0.93	0.81
Transformer	91.18%	0.90	0.96	0.85

Table 8: Scores on test data for binary classification.

and F1 score on our class 1 (i.e., ‘preclinical stage AD’). As can be seen, the transformer model outperforms others in terms of all the measured parameters. It is worth pointing out how measuring just accuracy could not always be a good metric when comparing models. In terms of accuracy, if we compare the baseline model with our transformer algorithm, the improvement is just 3.21%. However, with F1 score e.g., the improvement is much higher (e.g., 8.4% higher F1 score). This is why, to have a clear picture of which model is performing better, other statistical measures, such as F1 score, precision and recall, have to be taken into consideration.

Discussion

It is important to note that a subset of individuals with biomarker evidence of preclinical AD will not progress to developed AD dementia during their lifetime (Rowe et al. 2010). Hence, a fraction of the ‘cognitively normal’ labeled/diagnosed participants in the OASIS 3 dataset have preclinical AD biomarker, even though their disease never progressed to clinically diagnosed dementia stage. These individuals may contribute to the false positive (misclassification to AD) rate of our evaluation results.

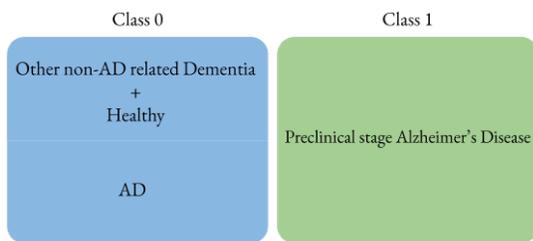


Figure 5: Overview of data labeling approach that includes AD dementia

In addition to using the labeling scheme shown in Fig. 3, we performed further experiments by following the data labeling scheme shown in Fig. 5. In this case, the ‘other’ class included patients with diagnosed AD, and we performed classification to differentiate individuals with ‘preclinical AD’ from others, including the patients with developed AD dementia. We compared the performance of 3D CNN and 3D RVN models. As shown in Table 9, the 3D RVN model achieves an F1 score of 0.78 and accuracy of 84.6%. Even though these results are not as high compared to the previously evaluated ‘preclinical AD’ assessment models (using the labeling scheme in Fig. 3), the glimpse locations identified by this RVN agent carry potential. More specifically, they highlight the brain regions responsible for differentiating ‘preclinical AD’ versus developed AD (stages 2 and 3).

These findings will enhance our understanding on the deterioration of the brain structure with progression of AD (from preclinical to AD stages 2 and 3).

	3D CNN	3D RVN
Accuracy	83.27%	84.62%
F1 Score	0.81	0.78
Precision	0.94	0.91
Recall	0.71	0.69

Table 9: Binary approach to differentiate preclinical AD versus others, including developed AD.

Conclusion

A lot of work has been done detecting MCI (stage 2), the earliest form of detectable Alzheimer’s disease (Markesbery 2010). The novelty and significance of our work is that we focus on and are able to predict AD at preclinical stage (stage 1), when all the parameters, physical assessments, and clinical data state that the patient is healthy.

With that goal, we adapted two different attention-based network models, and compared their performances with a 3D CNN-based baseline that we implemented. The baseline model uses stacked MRI scans for classification. The first attention model is a recurrent attention model, which extracts glimpses from stacked images, and feeds them into recurrent attention units to get a classification result. The second model is a modified and re-purposed transformer network, which first extracts features of a sequence of images from a pre-trained network, and then feeds these features to a transformer structure to be able to classify the sequence of images. Among these three approaches, the transformer model outperformed others achieving an F1 score of 0.90 and accuracy of 91.18% in preclinical (i.e., prior to MCI and dementia stage) Alzheimer disease detection. Experimental results demonstrate that, by using the MRI images effectively, it is possible to detect preclinical stage Alzheimer’s disease with a very promising accuracy.

References

- Ando, S.; and Huang, C. Y. 2017. Deep Over-sampling Framework for Classifying Imbalanced Data. volume 10534, 770–785. ISBN 0302-9743.
- Bauer, C. M.; Cabral, H. J.; and Killiany, R. J. 2018. Multimodal Discrimination between Normal Aging, Mild Cognitive Impairment and Alzheimer’s Disease and Prediction of Cognitive Decline. *Diagnostics (Basel)* 8(1): 14.
- Böhle, M.; Eitel, F.; Weygandt, M.; and Ritter, K. 2019. Layer-Wise Relevance Propagation for Explaining Deep Neural Network Decisions in MRI-Based Alzheimer’s Disease Classification. *Frontiers in aging neuroscience* 11: 194.
- Cheng, D.; Liu, M.; Fu, J.; and Wang, Y. 2017. Classification of MR brain images by combination of multi-CNNs for AD diagnosis. volume 10420, 875 – 879. International Society for Optics and Photonics, SPIE. doi:10.1117/12.2281808. URL <https://doi.org/10.1117/12.2281808>.

- Crous-Bou, M.; Minguillón, C.; Gramunt, N.; and Molinuevo, J. L. 2017. Alzheimer's disease prevention: from risk factors to early intervention. *Alzheimer's Research & Therapy* 9(1): 71.
- Ertekin, T.; Acer, N.; Köseoğlu, E.; Zararsız, G.; Sönmez, A.; Gümüş, K.; and Kurtoğlu, E. 2016. Total intracranial and lateral ventricle volumes measurement in Alzheimer's disease: A methodological study. *Journal of Clinical Neuroscience* 34: 133–139.
- Girdhar, R.; Carreira, J.; Doersch, C.; and Zisserman, A. 2019. Video Action Transformer Network. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* 244–253.
- Golde, T. E.; DeKosky, S. T.; and Galasko, D. 2018. Alzheimer's disease: The right drug, the right time. *Science* 362(6420): 1250–1251.
- Gupta; Yubraj; Lama; Kumar, R.; Kwon; and Goo-Rak. 2019. Prediction and Classification of Alzheimer's Disease Based on Combined Features From Apolipoprotein-E Genotype, Cerebrospinal Fluid, MR, and FDG-PET Imaging Biomarkers. *Frontiers in Computational Neuroscience* 13: 72. ISSN 1662-5188. doi:10.3389/fncom.2019.00072. URL <https://www.frontiersin.org/article/10.3389/fncom.2019.00072>.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Comput.* 9(8): 1735–1780. ISSN 0899-7667. doi:10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Hurd, M. D.; Martorell, P.; Delavande, A.; Mullen, K. J.; and Langa, K. M. 2013. Monetary costs of dementia in the United States. *New England Journal of Medicine* 368(14): 1326–1334.
- Islam, J.; and Zhang, Y. 2017. A Novel Deep Learning Based Multi-class Classification Method for Alzheimer's Disease Detection Using Brain MRI Data. volume 10654, 213–222. ISBN 0302-9743.
- Kang, M.; Kim, S.; Han, S.; Bae, J.; Yu, W.; Park, M.; Ku, S.; and Yang, Y. 2019. Effect of Paper-Based Cognitive Training in Early Stage of Alzheimer's Dementia. *Dementia and Neurocognitive Disorders* 18: 62. doi:10.12779/dnd.2019.18.2.62.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980.
- Korolev, S.; Safiullin, A.; Belyaev, M.; and Dodonova, Y. 2017. Residual and plain convolutional neural networks for 3D brain MRI classification. 835–838.
- Korolev, S.; Safiullin, A.; Belyaev, M.; and Dodonova, Y. 2017. Residual and plain convolutional neural networks for 3D brain MRI classification. *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)* 103(3): 200–204.
- LaMontagne, P. J.; Benzinger, T. L.; Morris, J. C.; Keefe, S.; Hornbeck, R.; Xiong, C.; Grant, E.; Hassenstab, J.; Moulder, K.; Vlassenko, A.; et al. 2019. OASIS-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer disease. *medRxiv*.
- Lee, G.; Nho, K.; Kang, B.; Sohn, K.-A.; Kim, D.; and for Alzheimer's Disease Neuroimaging Initiative. 2019. Predicting Alzheimer's disease progression using multi-modal deep learning approach. *Scientific reports* 9(1): 1952–12.
- Logsdon, R. G.; Gibbons, L. E.; McCurry, S. M.; Teri, L.; et al. 1999. Quality of life in Alzheimer's disease: patient and caregiver reports. *Journal of Mental Health and Aging* 5: 21–32.
- Lohrke, J.; Frenzel, T.; Endrikat, J.; Alves, F. C.; Grist, T. M.; Law, M.; Lee, J. M.; Leiner, T.; Li, K.-C.; Nikolaou, K.; et al. 2016. 25 years of contrast-enhanced MRI: developments, current challenges and future perspectives. *Advances in therapy* 33(1): 1–28.
- Manzano, S.; González, J.; Marcos, A.; Payno, M.; Villanueva, C.; and Matías-Guiu, J. 2009. [Experimental models in Alzheimer's disease]. *Neurologia (Barcelona, Spain)* 24(4): 255–262. ISSN 0213-4853. URL <http://europepmc.org/abstract/MED/19603296>.
- Markesbery, W. R. 2010. Neuropathologic Alterations in Mild Cognitive Impairment: A Review. *Journal of Alzheimer's Disease* 19: 221–8. ISSN 1387-2877. doi: 10.3233/JAD-2010-1220. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2872776/>.
- Park, J. S.; Chung, M. S.; Park, H. S.; Shin, D. S.; Har, D.-H.; Cho, Z.-H.; Kim, Y.-B.; Han, J.-Y.; and Chi, J.-G. 2010. A proposal of new reference system for the standard axial, sagittal, coronal planes of brain based on the serially-sectioned images. *Journal of Korean medical science* 25(1): 135–141.
- Perneckzy, R. 2018. *Biomarkers for preclinical Alzheimer's disease*. Springer.
- Reiman, E. M.; Quiroz, Y. T.; Fleisher, A. S.; Chen, K.; Velez-Pardo, C.; Jimenez-Del-Rio, M.; Fagan, A. M.; Shah, A. R.; Alvarez, S.; Arbelaez, A.; Giraldo, M.; Acosta-Baena, N.; Sperling, R. A.; Dickerson, B.; Stern, C. E.; Tirado, V.; Munoz, C.; Reiman, R. A.; Huentelman, M. J.; Alexander, G. E.; Langbaum, J. B.; Kosik, K. S.; Tariot, P. N.; and Lopera, F. 2012. Brain imaging and fluid biomarker analysis in young adults at genetic risk for autosomal dominant Alzheimer's disease in the presenilin 1 E280A kindred: a case-control study. *Lancet neurology* 11(12): 1048–1056.
- Rice, D. P.; Fillit, H. M.; Max, W.; Knopman, D. S.; Lloyd, J. R.; Dutttagupta, S.; et al. 2001. Prevalence, costs, and treatment of Alzheimer's disease and related dementia: a managed care perspective. *American Journal of Managed Care* 7(8): 809–820.
- Robbins, H. 2007. A Stochastic Approximation Method. *Annals of Mathematical Statistics* 22: 400–407.
- Rowe, C. C.; Ellis, K. A.; Rimajova, M.; Bourgeat, P.; Pike, K. E.; Jones, G.; Fripp, J.; Tochon-Danguy, H.; Morandau, L.; O'Keefe, G.; et al. 2010. Amyloid imaging results from the Australian Imaging, Biomarkers and Lifestyle (AIBL) study of aging. *Neurobiology of aging* 31(8): 1275–1283.

Schindler, S. E.; Jasielec, M. S.; Weng, H.; Hassenstab, J. J.; Grober, E.; McCue, L. M.; Morris, J. C.; Holtzman, D. M.; Xiong, C.; and Fagan, A. M. 2017. Neuropsychological measures that detect early impairment and decline in pre-clinical Alzheimer disease. *Neurobiology of aging* 56: 25–32.

Schlemper, J.; Oktay, O.; Schaap, M.; Heinrich, M.; Kainz, B.; Glocker, B.; and Rueckert, D. 2019. Attention gated networks: Learning to leverage salient regions in medical images. *Medical Image Analysis* 53: 197 – 207. ISSN 1361-8415. doi:<https://doi.org/10.1016/j.media.2019.01.012>. URL <http://www.sciencedirect.com/science/article/pii/S1361841518306133>.

Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556.

Sperling, R. A.; Jack, C. R.; and Aisen, P. S. 2011. Testing the right target and right drug at the right stage. *Science translational medicine* 3(111): 111cm33–111cm33.

Sutskever, I.; Martens, J.; Dahl, G.; and Hinton, G. 2013. On the importance of initialization and momentum in deep learning. volume 28 of *Proceedings of Machine Learning Research*, 1139–1147. Atlanta, Georgia, USA: PMLR. URL <http://proceedings.mlr.press/v28/sutskever13.html>.

Szegedy, C.; Ioffe, S.; Vanhoucke, V.; and Alemi, A. A. 2017. Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. *AAAI'17*, 4278–4284. AAAI Press.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30, 5998–6008. Curran Associates, Inc. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.

vital statistics system, N. 2019. Dementia Mortality in the United States, 2000–2017. https://www.cdc.gov/nchs/data/nvsr/nvsr68/nvsr68_02-508.pdf. Accessed: 09/15/2020.

Wood, D.; Cole, J. H.; and Booth, T. C. 2019. NEURO-DRAM: a 3D recurrent visual attention model for interpretable neuroimaging classification. *CoRR* abs/1910.04721.

Yu, F.; Rose, K. M.; Burgener, S. C.; Cunningham, C.; Buettner, L. L.; Beattie, E.; Bossen, A. L.; Buckwalter, K. C.; Fick, D. M.; Fitzsimmons, S.; Kolanowski, A.; Janet, K.; Specht, P.; Richeson, N. E.; Testad, I.; and McKenzie, S. E. 2009. Cognitive training for early-stage Alzheimer's disease and dementia. *Journal of gerontological nursing* 35(3): 23–29.

Zhou, D.-X. 2020. Theory of deep convolutional neural networks: Downsampling. *Neural networks* 124: 319–327.

Zott, B.; Busche, M. A.; Sperling, R. A.; and Konnerth, A. 2018. What happens with the circuit in Alzheimer's disease in mice and humans? *Annual review of neuroscience* 41: 277–297.