

Early Detection of Alzheimer's Disease From Cortical and Hippocampal Local Field Potentials Using an Ensembled Machine Learning Model

Marcos Fabietti¹, Member, IEEE, Mufti Mahmud², Senior Member, IEEE, Ahmad Lotfi³, Senior Member, IEEE, Alessandro Leparulo⁴, Roberto Fontana⁵, Stefano Vassanelli⁶, and Cristina Fasolato⁷

Abstract—Early diagnosis of Alzheimer's disease (AD) is a very challenging problem and has been attempted through data-driven methods in recent years. However, considering the inherent complexity in decoding higher cognitive functions from spontaneous neuronal signals, these data-driven methods benefit from the incorporation of multimodal data. This work proposes an ensembled machine learning model with explainability (EXML) to detect subtle patterns in cortical and hippocampal local field potential signals (LFPs) that can be considered as a potential marker for AD in the early stage of the disease. The LFPs acquired from healthy and two types of AD animal models ($n = 10$ each) using linear multielectrode probes were endorsed by electrocardiogram and respiration signals for their veracity. Feature sets were generated from LFPs in temporal, spatial and spectral domains and were fed into selected machine-learning models for each domain. Using late fusion, the EXML model achieved an overall accuracy of 99.4%. This provided insights into the amyloid plaque deposition process as early as 3 months of the disease onset by identifying the subtle patterns in the network activities. Lastly, the individual and ensemble models were found to be robust when evaluated by randomly masking channels to mimic the presence of artefacts.

Index Terms—Deep learning, dementia, neuronal signals, neuronal network, multimodal.

I. INTRODUCTION

ALZHEIMER'S disease (AD) is a neuropathology, which affects 46.8 million people worldwide [1]. While the disease is not fully understood, it is known that the lesions caused by amyloid plaques and tangles disrupt the connectivity of neurons, leading to their death and the atrophy of the brain. The main consequence is dementia, which impacts the person's ability to think, behave, work and function independently. The total healthcare cost for the treatment of AD in 2020 is estimated at US\$305 billion in the United States alone, with the cost expected to increase to more than US\$1 trillion as the population ages [2].

Improving the understanding of the disease is key for an early diagnosis and treatment, so the search for alternative and objective biomarkers is crucial. To this end, in recent years, data-driven solutions have been developed in order to tackle this disease [3]. There have been several studies to detect AD from neuroimaging and clinical data [4], [5], [6], [7], [8], [9] and also to manage patients who suffer from the disease [10], [11]. A data-driven strategy is based on collecting and analysing data to ultimately obtain new insights based on them. Among the methods for data analysis, machine learning (ML) techniques have become popular in recent years due to their capacity to deal with vast amounts of data and their ability to make accurate predictions based on them. The benefits of ML-based solutions include that they may be able to recognise new markers and can aid the physicians' decision-making, as early symptoms may be mistaken with other pathologies such as vascular dementia. These techniques gather the information presented to them to construct a model which can be used to make inferences about unseen data and have been widely used in diverse fields. Among the many ML methods, deep neural networks stand out. Their design was inspired by the biological counterpart, and they allow for the non-linear processing of information [12].

These models have been applied to the different technologies used to extract information about the brain's health. These range from medical images, data generated by Omics

Manuscript received 3 December 2022; revised 29 March 2023; accepted 26 April 2023. Date of publication 22 June 2023; date of current version 6 July 2023. Marcos Fabietti was supported by the Nottingham Trent University Vice Chancellor's Doctoral Fellowship. Mufti Mahmud was supported by Nottingham Trent University's Strategic Research Theme Springboard Fund. The data collection was supported by the Italian Ministry of Education, Universities and Research (MIUR) through PRIN-20175C22WM grant to Cristina Fasolato. (Corresponding author: Mufti Mahmud.)

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Animal Care Committee of the University of Padua and the Italian Ministry of Health under Application No. 522/2018-PR.

Marcos Fabietti and Ahmad Lotfi are with the Department of Computer Science, Nottingham Trent University, NG11 8NS Nottingham, U.K.

Mufti Mahmud is with the Department of Computer Science, Nottingham Trent University, NG11 8NS Nottingham, U.K., and also with the Medical Technologies Innovation Facility and the Computing and Informatics Research Centre of Nottingham Trent University, NG11 8NS Nottingham, U.K. (e-mail: mufti.mahmud@ntu.ac.uk; muftimahmud@gmail.com).

Alessandro Leparulo, Stefano Vassanelli, and Cristina Fasolato are with the Department of Biomedical Sciences, University of Padova, 35127 Padova, Italy.

Roberto Fontana is with the Department of Physiology and Pharmacology, Sapienza Università di Roma, 00185 Rome, Italy.

Digital Object Identifier 10.1109/TNSRE.2023.3288835

technologies [13], [14], and neural signals, including electroencephalography (EEG), magnetoencephalography (MEG), electrocorticography (ECoG), local field potentials (LFP) and neuronal spikes [15]. Imaging modalities are the main application for machine learning classification of AD [16], which can be attributed to easier access to data, more explainability and the facility of using models developed for images such as convolutional neural networks (CNNs) [17]. The integration of multimodal data can help improve the diagnosis by providing complementary information about the subject's health [18].

Comparatively, neuronal signals allow studying the brain's activity over time, showing how different structures respond to a given task. By using invasive recordings, the activity of specific structures or populations of neurons can be studied, revealing new information thanks to probes with higher spatial resolution. We find the majority of work to be oriented towards EEG and MEG in humans [19]. This is due to the fact that both are non-invasive recordings, which facilitates the signal acquisition process. However, if a deeper structure, such as the hippocampus, is to be studied, the EEG/ECoG/MEG are limited as they only capture cortical information. When recording neuronal signals, other electrical sources, either instrumental or physiological, may distort the signals. They are commonly known as artefacts, and their identification and removal are of importance in order to further analyse and infer insights from them. One of the most common approaches is to discard the affected epochs, however, it may cause distortions in the output of a classification model [20]. Therefore, it is important to establish the stability of a model if the information is missing.

As research subjects, transgenic mouse lines are accepted as AD amyloidosis research models [21] and commonly employed to investigate medicines capable of altering illness progression and restoring memory functions [22]. Beker et al. researched the use of ML to identify AD in animal models, where a support vector machine model was applied in classifying LFPs into wild type (WT) or AD model (B6C3 APP^{swe}/PS1^{dE9}) and achieved an accuracy of 82.6% [23]. The model was fit with different features such as correlation coefficients (R^2), fit sum squared error, trough frequency, trough amplitude standard deviation, half-width amplitude, AC maximal power, Gaussian fit separation, coefficient of variation, inter spike interval, frequency maximum power, mean signal to noise ratio and amplitude kurtosis. The most discriminative feature for the LFPs was the R^2 , as normal neuronal activity is bimodal. In contrast, the AD model has a weaker bimodality due to struggles between regular state transitions [23].

When using an ML-based approach in neuronal signal classification, features from the different domains, such as temporal, spatial and spectral, can be extracted. In order to exploit the features of each domain, we trained and compared different ML models for each one: deep neural networks for the temporal domain, CNNs for the spatial domain and aggregated decision trees for the spectral domain. These were chosen for their renowned performances as state-of-the-art in a diverse range of applications.

Whilst these ML models can give accurate predictions, they can show no explanation about the decision-making process that leads to them. In order to develop models grounded in

reality, which follow our understanding of the effects of the disease, explainability is required. The absence of explainability limits their adoption by healthcare professionals due to a lack of "trust" in a black-box model, and it fails to address questions of potential biases. ML models should be in-the-loop with the end users, resulting in reduced errors and mitigating risks, more confidence and adoption of the system, improved model performance via user feedback, and, most importantly, delivering accurate data-driven decision-making.

In this study, we utilised animal models (AD and healthy) to record brain activities as LFPs and applied appropriate ML models in order to accurately detect the disease at its early stage. The decision-making process of the models was also explored to map them to known AD markers using activation maps. Late fusion of features was employed for the best-performing models from each domain in order to improve the overall detection performance. Furthermore, the robustness of the models to artefacts was evaluated by randomly masking channels to simulate the removal of affected segments and the classification of stress-period segments. To summarise, the main contributions of this study are:

- Identify the most discriminative features across the three domains and the most effective models for them.
- Develop and validate an ensembled ML model with explainability (EXML) using known AD biomarkers.
- Develop a robust late fusion technique for early detection of AD in invasive neuronal signals.

The remainder of the article is partitioned into the methodology in Section II. In Section III we showcase the results obtained by the different models for each feature set and discuss them. Lastly, Section IV concludes the article.

II. METHODOLOGY

A block diagram of the proposed EXML model is shown in Fig. 1. Subsequent sections describe the model in detail.

A. Data Acquisition

For the analysis, three respective sets of ten three-month-old female healthy C57BL/6J (WT), single transgenic PS2.30H (ST) and double transgenic B6.152H (DT) mice were used, totalling thirty study subjects. All of the animals were kept in a specific pathogen-free animal facility with a 12-hour duration of light and dark cycles, and unrestricted access to food and water. All experimental procedures were performed according to the European Committee guidelines (decree 2010/63/CEE) and the Animal Welfare Act (7 USC 2131), in compliance with the ARRIVE guidelines, and were approved by the Animal Care Committee of the University of Padua and the Italian Ministry of Health (authorisation decree 522/2018-PR). For more details, we refer the reader to the published experimental analysis [24].

The mice were sedated with an intraperitoneal injection of urethane dissolved in 0.9 per cent NaCl and a combination of xylazine/tiletamine-zolazepam mixed in phosphate buffer. The spontaneous LFP activity was recorded using a linear 32-electrode silicon probe coupled to the acquisition equipment through a 32-channel head stage and a serial peripheral interface connection.

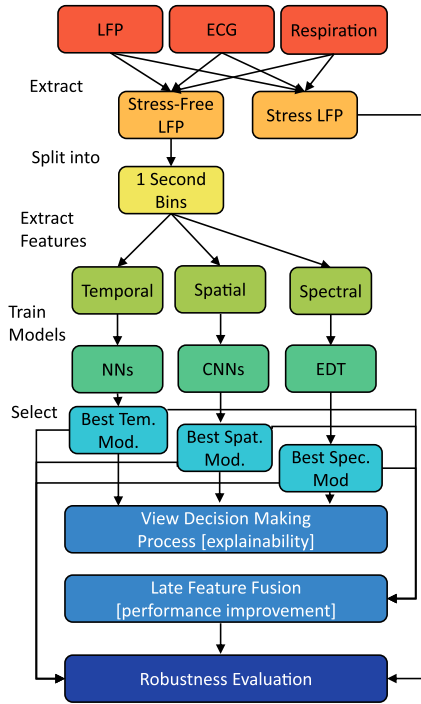


Fig. 1. The proposed ensemble ML model. Individual models considered are neural networks (NNs) for the temporal domain, convolutional neural networks (CNNs) for the spatial domain and ensemble decision trees (EDT) for the spectral domain.

Only 24 channels of the probe were used for recording, as they cover the area of interest from the first layer of the cortex to the deepest layer of the α dentate gyrus (DG) of the hippocampus, and the last 8 of the rest are out of brain tissue. This was chosen since AD mainly affects DG, CA1 and cortical layers [24], as well as to avoid unnecessary effects on deeper nuclei, thus the probe was introduced into the posterior parietal cortex and lowered to 2.3 mm.

Through the open graphic user interface software included with the Open-Ephys acquisition equipment, the LFPs were visualised, captured, and digitalised at 10 kHz. Other physiological signals were captured at the same time as the LFPs to evaluate the animal's health status. Electrocardiogram (ECG) recordings were 10X amplified and filtered between 1 and 100 Hz using a DAM 50 Amplifier to monitor the heartbeat. Positive and negative electrocardiogram derivations were placed subcutaneously into the forelimbs. The piezoelectric properties of the temperature probe converted respiration-induced chest wall motions into voltage variations. A DP-301 amplifier was used to amplify the respiration signal 100X and band-pass it between 0.1 and 100 Hz. Physiological signals were digitalised at 10 kHz using a PCI-6071E I/O card in a differential mode in conjunction with a BNC-2090 terminal block and recorded using a custom-written LabView script. After obtaining a stable level of anaesthesia, spontaneous brain activity was recorded for 30-40 minutes.

B. Pre-Processing

Offline processing of the LFPs was carried out in Matlab utilising custom-written scripts. At first, the raw signal files

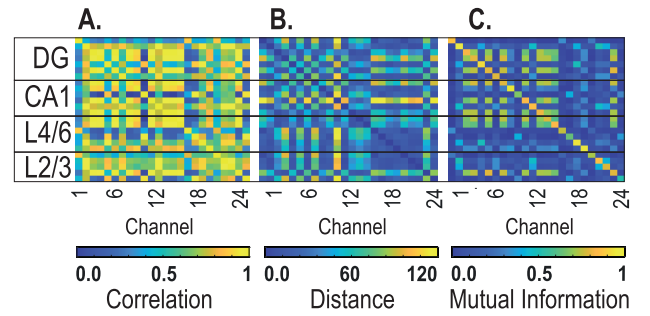


Fig. 2. Spatial feature maps: correlation (A), dynamic time warping (B) and mutual information (C).

are converted to Matlab files from Open-Ephys format. Subsequently, the 50 Hz line noise is removed by applying a Gaussian filter. The first 24 channels' signals were filtered using the built-in non-causal zero-phase distortion filtering algorithm, which in order to avoid phase distortion, the data is processed in both forward and reverse directions using coefficients from the built-in Butterworth transfer function. Using a median estimation approach, baseline drift was eliminated from all signals (LFP, ECG, and respiration). Afterwards, the recordings were low-pass filtered (filter order: 5; cut-off frequency: 190 Hz for LFP, 25 Hz for ECG, and 10 Hz for respiration) and down-sampled to 500 Hz, 50 Hz and 20 Hz, respectively. The supporting respiratory and ECG signals were used to find time periods when a rodent was feeling well. From the longest segment without stress, non-overlapping one-second windows across all channels were extracted from each rodent. As there are fewer normal state feature segments from ST mice, both WT and DT were downsampled to match the number of ST examples. For cross-validation we chose the hold-out method, where the final dataset consisting of 20,328 balanced examples was split 70%-15%-15% for training, validation, and testing, respectively.

From the examples, different features were extracted for each of the models. For the temporal models, the raw signal was used. For the spatial model, different metrics were used to evaluate the similarities between two given channels. First, a correlation matrix was built based on the Pearson correlation of a channel x and y , given the equation:

$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2 (y_i - \bar{y})^2}} \quad (1)$$

where $i = 1, \dots, 500$. For comparison, a dynamic time warping (DTW) algorithm was used [25].

In addition, spatial maps based on the mutual information (MI) [26] of any two channels x and y were extracted using the following equation:

$$MI(x; y) = \sum_{i=1}^n \sum_{j=1}^n p(x(i), y(j)) \cdot \log \left(\frac{p(x(i), y(j))}{p(x(i)) \cdot p(y(j))} \right) \quad (2)$$

where $i, j = 1, \dots, 500$. Examples of the three extracted spatial features of a WT example are shown in Fig. 2.

Lastly, for the spectral models, the spectral bands were defined as slow oscillations (0.1-1.7 Hz), delta (1.7-4.7), theta (4.7-10), beta (10-25 Hz), low gamma (25-45 Hz), high gamma (45-90 Hz) and fast oscillations (90-125 Hz). First, the signals across the channels were averaged in the selected areas (1-24 all, 1-7 DG, 8-13 CA1, 14-19 cortex layers 4 to 6 (L4/6), 20-24 cortex layers 2 to 3 (L2/3)). Subsequently, the absolute and relative power of all the bands were extracted, and the ratio of the power of slow oscillation to delta bands and the ratio of the power of the bands from 0.1-4.7 Hz to 4.7-125 Hz were also used.

C. Machine Learning Models

1) *Models for Temporal Domain*: For temporal models, the raw signal is used as an input. That means that the method must be able to extract relevant features by itself, which is where neural networks shine. Inspired by the brain's working principle, the neural networks are composed of multiple layers of processing of non-linear information to analyse patterns and classify them. Since their development, there have been many improvements in architecture and in their complexity. For the temporal models, we first explore the traditional multi-layered perceptron (MLP), followed by a state-of-the-art CNN for EEG recordings EEGNET [27] (imported from keras-python). Subsequently, we employ models known to perform well for time series, that is a recurrent (i.e. long-short-term memory (LSTM)) network that averages the inputs across channels in order to feed it to said layer. Lastly, a combination of LSTM with CNN as a feature extractor, the LSTM-CNN network. The layers of these networks are illustrated in Fig. 3(A), and the operations carried out by each one are as follows:

a) *Fully connected*: The output of fully connected layer y_i given an input x_i and W the matrix of weights, can be expressed as:

$$y_i = f(Wx_i + b_i) \quad (3)$$

b) *Dropout*: The dropout layer y_i turns off a random set of neurons, expressed as:

$$y_i = (1 - p)f(Wx_i + b_i) \quad (4)$$

where p is the dropout distribution.

c) *LSTM*: The operations of a cell within an LSTM are:

$$\begin{aligned} f_i &= \sigma(W_{fh}h_{i-1} + W_{fx}x_i + b_f), \\ u_i &= \sigma(W_{ih}h_{i-1} + W_{ix}x_i + b_i), \\ \tilde{c}_i &= \tanh(W_{\tilde{c}h}h_{i-1} + W_{\tilde{c}x}x_i + b_{\tilde{c}}), \\ c_i &= f_i \cdot c_{i-1} + u_i \cdot \tilde{c}_i, \\ o_i &= \sigma(W_{oh}h_{i-1} + W_{ox}x_i + b_o), \\ h_i &= o_i \cdot \tanh(c_i) \end{aligned} \quad (5)$$

where the variable x_i is the input vector, W are the weights, b is the bias and σ is the sigmoid function. Additionally, f_i is the forget gate, u_i is the update gate, \tilde{c}_i is the cell input, c_i is the cell state, o_i is the output gate and h_i the hidden state or output vector of the cell at time i .

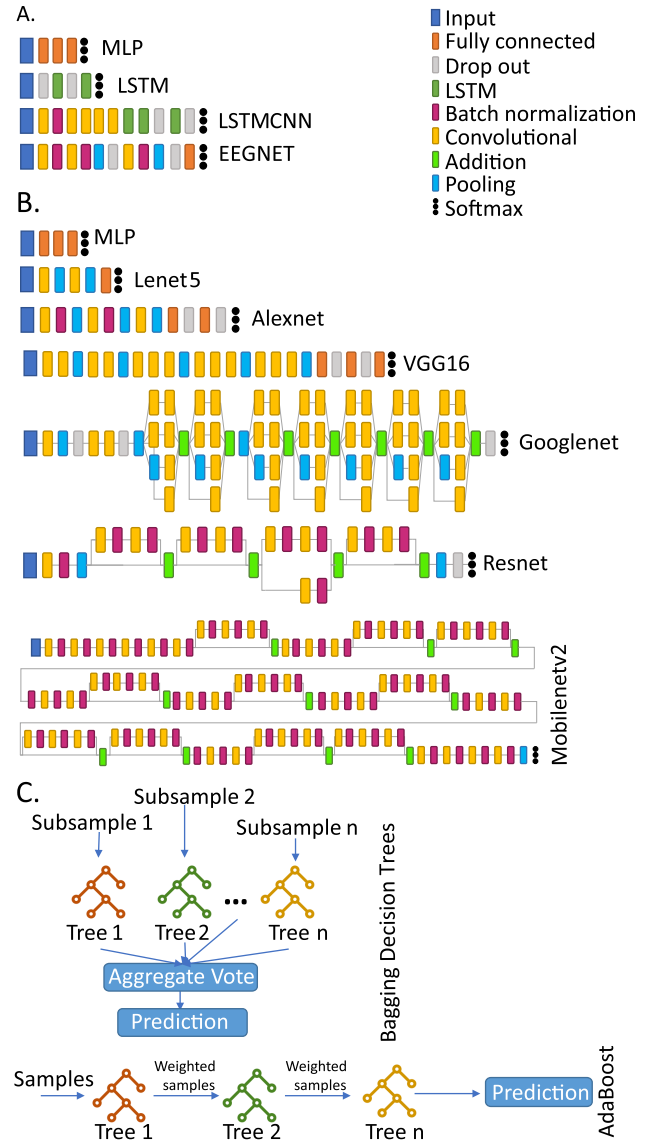


Fig. 3. Architectures of the temporal models (A): MLP, LSTM, LSTM CNN and EEGNET; the spatial models (B): MLP, Lenet-5, Alexnet, VGG16, Googlenet, Resnet and Mobilenetv2; and the spectral models (C) Bagged decision trees and AdaBoost.

d) *Batch normalisation*: Given input vectors x_i , batch normalisation output y_i is expressed as:

$$y_i = \gamma \left(\frac{x_i - \mu_B}{\sqrt{\sigma_B^2 + \epsilon}} \right) + \beta \quad (6)$$

where μ_B mini-batch mean, σ_B^2 mini-batch variance, γ scaling and β shift.

e) *Convolution Layer*: Given an input vector $x_{i_{hw}}$ of $m \times n$ dimensions where $h = 1, \dots, m$ and $w = 1, \dots, n$, and filter H of size $l \times l$, the resulting vector $y_{i_{hw}}$ of the convolution between the two is expressed as:

$$y_{i_{hw}} = \sum_{u=1}^l \sum_{s=1}^l H_{us} x_{i_{(h,u)(w,s)}} \quad (7)$$

f) *Addition*: An addition layer adds inputs from multiple neural network layers element-wise. For given two inputs x_{i_1} and x_{i_2} , the output y_i is defined as:

$$y_i = x_{i_1} ++ x_{i_2} \quad (8)$$

g) *Pooling*: For a feature map x_i having dimensions $x_{i_h} \cdot x_{i_w} \cdot x_{i_c}$ the dimensions of output y_i obtained after a pooling layer are:

$$\frac{(x_{i_h} - f + 1)}{s} \cdot \frac{(x_{i_w} - f + 1)}{s} \cdot x_{i_c} \quad (9)$$

where s is the stride and f the filter size.

h) *Softmax*: The activation function softmax is used to calculate d class probabilities, expressed as:

$$y_i = \frac{\exp(x_i)}{\sum_{j=1}^d \exp(x_j)} \quad (10)$$

where the label \hat{y} is defined as the maximum value $\text{argmax}(y_i)$. For the temporal networks, the optimisation algorithm used was Adam, with an initial learning rate of 0.0001, a momentum of 0.9 and a batch size of 128.

2) *Models for Spatial Domain*: Brain connectivity is affected by AD [28], thus, it is important to look at the relationship between different structures through spatial maps. The different extracted spatial maps can be considered images of 24 by 24 pixels. For this task, convolutional neural networks are the prime candidate. CNNs are a specific form of neural network that is well suited to computer vision applications due to their capacity to hierarchically abstract representations of spatial operations.

Adaptations of different popular architectures such as Lenet-5, Alexnet, VGG16, Googlenet, Mobilenetv2 and Resnet were explored [29]. Networks were adjusted in order to work with an input map of 24×24 , including reducing the number of filters and filter size, reducing the fully connected layer's neurons, and in those where the repeated number of pooling layers limit the input (i.e. Googlenet and Resnet), iterative modules were removed. In order to understand the impact of network connectivity, Grad-CAM heatmaps were extracted and averaged for each class of the test set. They are generated by visualising the gradients of the last convolution layer's feature maps for a given signal. The visualisations are class-discriminative; in other words, a feature area is associated with the class with the strongest reliance on that area. The layers of these networks are described in Fig. 3(B). The optimisation algorithm used was Adam, with an initial learning rate of 0.001, a momentum of 0.9 and a batch size of 256 for all models.

3) *Models for Spectral Domain*: Unlike the temporal model, which took time series as inputs, and the spatial models, which took images as inputs, for the spectral models the input are individual features. While there is a large range of methods available for this task, we are looking for explainable models. Classification and regression trees [30] can be used to build both classifications and regression trees, where each internal node has exactly two outgoing edges, namely binary trees. The splits are selected using the Gini index as a splitting criterion

and the obtained tree is pruned by cost complexity pruning. The Gini index is described in the following equation:

$$Gini = 1 - \sum_{i=1}^n (p_i)^2 \quad (11)$$

where p_i is the probability of an object being classified to a particular class. Binary decision trees are easier to interpret than other ML models, as the decision process in the tree is easily apparent, and it mimics the rule-based systems used for medical diagnosis. Decision trees are machine learning methods in which the decision process is both local and global and integrated into the methods themselves.

The performance of decision trees can be further improved by aggregating models, namely bagging and boosting [31]. For bagging, given a dataset $D = (x_1, y_1), \dots, (x_n, y_n)$, sample T sets of n elements from D (with replacement), obtain $D_1, D_2, \dots, D_T \rightarrow T$ quasi replica training sets, followed by training a machine on each $D_i, i = 1, \dots, T$ and obtain a sequence of T outputs $f_1(x), \dots, f_T(x)$. The final classification score is:

$$\bar{f}(x) = \text{sign}\left(\sum_{i=1}^T f_i(x)\right) \quad (12)$$

Ad-hoc analysis indicated that the bagged trees method of utilising all features in each split showed better performance than utilising a subset of them, a variation known popularly as random forests, thus the former was chosen. For boosting we utilise the multi-class version of Adaboost, AdaBoost.M2 [32]. We use these ensemble classifiers for the spectral models, and based on the best-performing one across the different areas we explored the performance of the classifier for different permutations of the classifications. Furthermore, the predictor importance of features can be calculated by summing changes in the node risk due to splits on every predictor and then dividing the sum by the total number of branch nodes. This allows for gaining insight into the selection of markers.

D. Ensemble Model and Late Fusion

In areas such as neuroimaging the fusion of different modalities to improve the understanding of pathologies and their diagnosis is widely used. We propose the use of feature fusion across the best-performing models in order to improve the classification results, as it has been successfully applied to non-invasive neuronal recordings [33]. As we want to maintain insight into the different modalities provided by the individual models, late fusion models are proposed.

Denoting P_m the confidence (or probability) score yielded by classification models $ML_m, (m = 1, \dots, n)$, where $n = 3$, ML_1 is the temporal model, ML_2 comes from the spatial model, and ML_3 refers to the spectral model. Four fusion rules are considered: average fusion, maximum fusion, minimum fusion, and Naive-product fusion. The average rule simply calculates the simple mean of the ML classifiers' outputs.

$$f_{mean} = \frac{1}{n} \sum_{m=1}^n P_m \quad (13)$$

The maximum rule outputs the maximum value over the classifier responses,

$$fmax = \max(P_m) \quad (14)$$

while the minimum rule is

$$fmin = \min(P_m) \quad (15)$$

Assuming classifiers' independence given the different domains, the Naive-product rule is expressed by.

$$fNprod = \frac{\prod_{m=1}^n P_m}{\prod_{m=1}^n P_m + \prod_{m=1}^n (1 - P_m)} \quad (16)$$

E. Robustness to Artefacted Channels

Lastly, the performance of each domain's model and the ensemble model is evaluated by increasingly removing random channels in order to find the optimal model. The process is described in Algorithm 1, where given a matrix $M[C, T, N]$ of the neuronal recording, $C = 1, \dots, 24$ are the total number of channels, $T = 1, \dots, 500$ are the datapoints, $N = 1, \dots, 20328$ are the number of examples, $L = 1, \dots, 12$ are the channels are selected for occlusion in each iteration, $randi(a,b,c)$ is the built-in Matlab random integer generation of a matrix of b by c numbers ranging from 1 to a , and the final result OcM is a 4D matrix. The selected channels are replaced by the average of the remaining ones, in order to maintain spatial cohesion.

Algorithm 1 Random Channel Occlusion Algorithm

Input: M, L

Output: OcM

Initialisation:

```

1:  $M[C, T, N]$  = sample matrix
2:  $C$  = total number of channels
3:  $T$  = data points
4:  $N$  = number of samples
5:  $L$  = maximum replaced channels,  $C/2$  in our case
6: for  $l := 1$  to  $L$  do
7:    $OcM[:, :, :, l] = M$ 
8:   for  $n := 1$  to  $N$  do
9:      $Oc = randi[C, l, 1]$ 
10:    remaining_ch =  $OcM[:, :, n, l]$ 
11:    remaining_ch[ $Oc, :, n, l$ ] = [ ];
12:    mean_remaining_ch = mean[remaining_ch,1]
13:    for  $o := 1$  to  $l$  do
14:       $OcM[OC[o], T, n, l] = mean\_remaining\_ch$ 
15:    end for
16:  end for
17: end for
18: return  $OcM$ 

```

III. RESULTS AND DISCUSSION

The performances of the models for each domain are shown in the subsequent sections.

TABLE I

PERFORMANCE OF TEMPORAL MODELS ON THE TEST SET

Model	accuracy	sen. (WT)	sen. (DT)	sen. (ST)
MLP	82.3	86.1	89.4	70.6
LSTM	59.2	60.9	86.1	29.0
LSTM-CNN	76.6	96.6	99.7	32.5
EEGNET	91.2	92.9	98.7	82.0

TABLE II

PERFORMANCE OF SPATIAL MODELS ON THE TEST SET

Function	Model	acc.	sen. (WT)	sen. (DT)	sen. (ST)
Corr	MLP	84.0	99.7	100.0	51.5
	LeNet	86.1	92.0	94.0	71.9
	AlexNet	97.3	99.8	100.0	91.9
	VGG16	98.3	98.8	98.1	98.8
	ResNet	87.5	99.7	99.9	62.4
	MobileNetv2	98.1	98.8	99.8	95.7
	GoogLeNet	99.1	99.0	98.9	99.4
MI	MLP	90.5	99.8	99.8	71.5
	LeNet	77.5	90	93.5	48.3
	AlexNet	85.2	98.4	99.5	57.2
	VGG16	70.8	79.9	96.3	35.7
	ResNet	84	99.4	99.3	52.5
	MobileNetv2	85	99.1	99.8	55.3
	GoogLeNet	98.2	97.7	97.9	99.1
DTW	MLP	85.6	99	99.9	57.1
	LeNet	88.8	98.3	99.2	68.4
	AlexNet	90.3	99.7	99.6	71.1
	VGG16	99.4	99.1	99.8	99.4
	ResNet	83.8	99.6	99.9	51.1
	MobileNetv2	88.3	99	99.6	65.8
	GoogLeNet	98.9	98.5	99.9	98.3

A. Performance of Machine Learning Models

1) *Models for Temporal Domain:* Results of the temporal models are shown in Table I. The MLP has been able to successfully classify all three classes, however, the LSTM and LSTM-CNN models have struggled to differentiate WT and ST from DT, leading to poorer performance. The best performance was achieved by the EEGNET, which can be attributed to the fact that the model not only exploits the temporal information but also the spatial properties due to depth-wise convolutions embedded in the architecture. However, it also struggles to differentiate between the waveforms of ST and the other rodents, as it resembles the DT.

2) *Models for Spatial Domain:* The compilation of performance of the neural networks in the test set is compiled in Table II. Of the correlation connectivity maps models, the MLP, LeNet, and ResNet achieved the worst performances due to the poor sensitivity of the ST. Oppositely, Alexnet, Mobilenetv2, VGG16, and GoogLenet achieved the best scores with accuracies over 97%. For MI connectivity maps models, most models were not able to accurately identify ST, averaging less than 50% sensitivity for Lenet, AlexNet, VGG16, ResNet, and MobileNetv2. On the other hand, MLP and Googlenet achieved scores of 71.5% and 99.6%, respectively. Lastly, for DTW connectivity maps models the best scores are achieved by GoogLenet, followed closely by VGG16. The remainder of the models have a lower accuracy due to a lower ST sensitivity, ranging from 51.1% of the ResNet up to 71.1% of the AlexNet.

Overall, correlation spatial maps perform better than the MI and the DTW spatial maps. While all networks achieved good sensitivity for the WT and DT classes, most struggled to

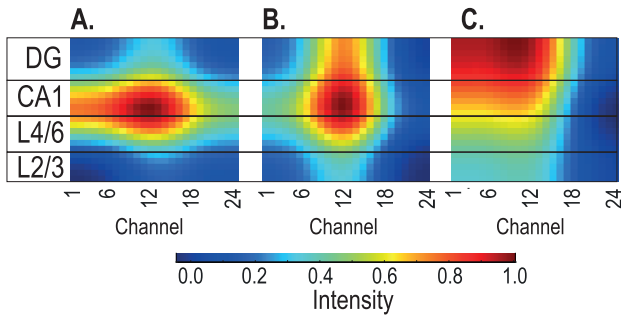


Fig. 4. Averaged GRADCAM map for WT (A), DT (B) and ST (C).

TABLE III

PERFORMANCE OF ENSEMBLE TREE PER AREA ON THE TEST SET

Ensemble	Area	accuracy	sen. (WT)	sen. (DT)	sen. (ST)
Bagged	DG	84.9	83.9	87.1	83.9
	CA1	81.4	80.4	82.8	81.1
	L4/6	85.4	83.8	86.1	86.4
	L2/3	81.1	81.2	82.1	79.9
	ALL	84.7	82.6	88.8	82.9
Adaboost	DG	75.7	69.0	77.0	81.0
	CA1	74.4	66.5	79.0	77.6
	L4/6	76.4	68.6	78.5	81.7
	L2/3	71.9	67.7	75.1	73.0
	ALL	75.7	69.9	79.4	77.7

distinguish ST, which was miss-classified as DT. Nonetheless, the GoogLenet model was the only model to consistently achieve good performance across all spatial maps. We proceed to use it to extract the average grad cam for each class in the test set, which is shown in Fig. 4. The GRAD CAMs (A) and (B) reveal that the model focuses on the local connectivity of the CA1 area for both WT and DT. Similarly, GRAD CAM (C) reveals that for the ST the CNN looks at the connectivity of DG with CA1. These areas coincide with what has been reported in the literature of where AD manifests in the earliest stages, for both mice model of AD [34] and in humans [35], since at this stage the connectivity between cortex and it has not been significantly affected so is not a marker for the network.

3) *Models for Spectral Domain*: The results achieved on the test set by the ensemble tree models per area are compiled in Table III. Overall, the bagged models outperform the Adaboost models. This indicates that a single tree possesses high variance and that by bagging we are improving the stability of the model, which suits the data better instead of improving the complexity of the model i.e. boosting. For the bagged trees, across the different areas, there isn't a major drop in the classifier's performance. The best sensitivity for WT is obtained in the DG, and the same with DT, while for the ST is at L4/6. In the case of Adaboost, the best sensitivity for WT and DT is across all while for the ST is at L4/6 as well.

With these results in mind, we explore further the different classification combinations for the bagged tree model, as shown in Table IV. The model works best in binary classifications, especially WT vs DT with accuracy over 94%, however, the performance drops 5-10% when doing a three-way classification. This means that the distribution of the

TABLE IV

PERFORMANCE OF BAGGED TREES FOR DIFFERENT CLASSIFICATIONS ON THE TEST SET

Classification	Area	acc.	sen. (WT)	sen. (DT)	sen. (ST)
WT VS ST VS DT	DG	84.9	83.9	87.1	83.9
	CA1	81.4	80.4	82.8	81.1
	L4/6	85.4	83.8	86.1	86.4
	L2/3	81.1	81.2	82.1	79.9
	ALL	84.7	82.6	88.8	82.9
WT VS DT	DG	94.3	92.8	95.8	
	CA1	91.3	90.5	92.0	
	L4/6	93.5	92.8	94.3	
	L2/3	92.2	89.9	94.5	
	ALL	94.2	93.3	95.1	
WT VS ST	DG	90.6	87.8		93.5
	CA1	89.3	88.0		90.7
	L4/6	90.8	88.0		93.6
	L2/3	88.9	85.1		92.7
	ALL	89.3	87.8		90.9
ST VS DT	DG	89.4		89.4	89.4
	CA1	88.8		88.1	89.5
	L4/6	88.6		88.9	88.3
	L2/3	86.1		85.9	86.3
	ALL	89.6		90.8	88.4

TABLE V

FEATURE RANKING PER AREA OF THE SPECTRAL MODEL

Feature	Ranking per Area				
	DG	CA1	L4/6	L2/3	ALL
fast oscillation	1	1	1	1	1
relative fast oscillation	2	2	2	2	3
high gamma	3	3	3	3	2
relative high gamma	7	4	7	8	8
low gamma	4	5	8	7	4
relative low gamma	8	6	5	5	7
beta	6	7	6	4	6
relative beta	5	8	4	6	5
theta	9	9	11	9	9
relative theta	13	13	13	12	13
delta	10	10	9	11	10
relative delta	11	11	10	10	11
slow oscillation	15	14	15	15	14
relative slow oscillations	16	16	16	16	16
slow oscillation delta ratio	14	15	14	14	15
low high ratio	12	12	12	13	12

biomarkers for differentiating the ST AD model may overlap more with the DT or WT distributions.

The feature importance of the different three-class bagged models is listed in Table V. The models give more importance to higher frequency features, such as absolute and relative fast oscillations, followed by absolute high and low gamma. On the other hand, low frequencies such as absolute and relative slow oscillations are ranked last. Both features which are ratios are also ranked low, indicating they are not markers for the ML model.

Neural network activity is aberrantly increased in AD patients and animal models due to functional deficits in and decreased activity of GABA inhibitory interneurons [36]. In particular, Kalemaki et al. [37] found that GABAergic inhibition produces significant reduction and disorganisation of the gamma frequency range (30-80 Hz) with an aberrant peak at high gamma frequency range (80-150 Hz), which translates to our high gamma (45-90 Hz) and fast oscillations (90-125 Hz) ranges. This correlates with the model finding

TABLE VI

PERFORMANCE OF THE FEATURE FUSION CLASSIFIERS AND THE INDIVIDUAL MODELS ON THE TEST SET

Model	accuracy	sen. (WT)	sen. (DT)	sen. (ST)
temporal	91.2	92.9	98.7	82.0
spatial	99.1	99	98.9	99.4
spectral	84.7	82.6	88.8	82.9
mean fusion	98.8	97.4	98.9	100
max fusion	99.3	98.8	99.2	100
min fusion	99.4	98.7	99.7	100
nav. prod. fusion	99.2	98.4	99.4	100

TABLE VII

STATE-OF-THE-ART COMPARISON

Author	Method	Classes	Accuracy
Beker et al. [23]	linear SVM	2	82.6
Proposed	Temporal	3	91.2
Proposed	Spatial	3	99.1
Proposed	Spectral	3	84.7
Proposed	Min Fusion	3	99.4

these as discriminative features. The role of these bands is that the gamma synchrony across the hippocampus plays a central role in the coordinated reactivation of stored memories [38].

In the literature, Arroyo-García et al. [39] reported that amyloid plaque accumulation in mouse model correlated with degradation of the gamma oscillations, which leads to consequent cognitive impairment. From a different angle, Swietlik et al. [40] simulated the DG-CA3-CA1 region of a pathological model, and found that inducing gamma oscillations could potentially minimise AD-related pathology. In human subjects, Van Deursen et al. [41] also found that gamma band power recorded via EEG was elevated in 15 AD patients compared to both mild cognitive impairment patients (20 subjects.) and the control group (20 subjects), proposing it could be used as a reliable biomarker. Overall, we find our results consistent with what is understood about AD.

B. Performance of Ensemble Model and Late Fusion

The different feature fusion models of the best models for each domain are listed in Table VI. All but the mean fusion improve the performance of the individual classifiers, specifically the sensitivities to the DT and the ST classes, which are improved to above 99%. Out of all of them, the best performing model is the min fusion. This means that if any of the three models is certain that an example does not belong to a said class, that is the most accurate decision. It is important to highlight that none of these ensemble methods requires training, unlike other approaches such as regressors, and achieves good results.

Looking at the individual models, the best-performing one is the spatial model, meaning it has the most discriminative features. This correlates with the fact that the best temporal model is EEGNET because it actually exploits spatial properties due to depth-wise convolutions embedded in the architecture. Lastly, the spectral model has a higher sensitivity to the ST than the temporal model, but a lesser overall performance.

In Table VII we compare our results with the method proposed by Beker et al. [23]. They proposed a binary

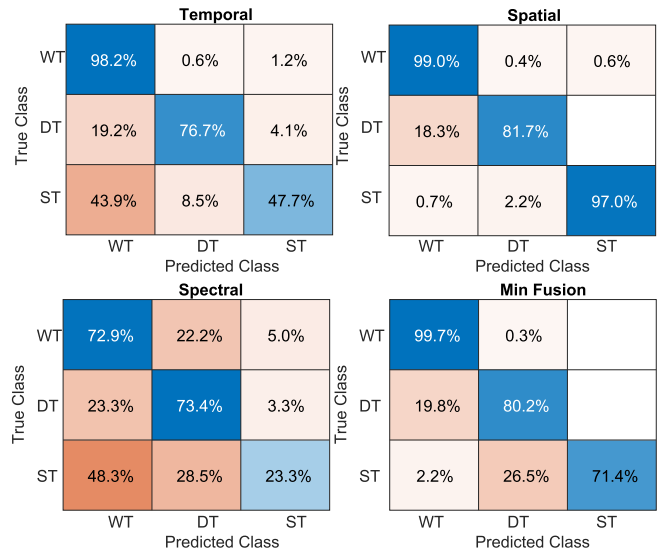


Fig. 5. Confusion matrices of the stress segments classified by the best models for each domain and the proposed ensemble.

classification model, detecting WT and the animal AD model, based on a linear support vector machine fed with hand-crafted features. Overall, all the proposed models show higher accuracy and are able to do so in a multi-class scenario. As shown in Table IV, the performance of these models would be improved further when looking at binary classification between the control and the AD model. Furthermore, we are able to extract information such as the relevant brain areas and spectral features that can be mapped to human subjects.

Having obtained the best models for each modality and ensemble, as well as comparing them against the known literature, their robustness when the rodent is under stress and when channels are missing due to artefacts is evaluated next.

C. Robustness to Stress Segments

The extracted stress segments for each rodent type were classified by the best-performing model of each domain and the proposed ensemble model. The results are illustrated in Figure 5 in the form of confusion matrices to highlight the type of incorrect classifications made by the models.

In the temporal model, the detection of WT maintains a high performance. On the other hand, the performance of the DT drops due to a miss-classification into WT and ST. Similarly, the ST performance achieves a near-guess level due to the model identifying it as WT. This was expected due to the abnormal waveforms of the signals in the stressed state. The spatial model correctly classifies most examples, as the stress does not significantly disturb the network's connectivity. However, there is a decline in the detection of DT.

The spectral model shows the lowest performance out of the three. The sensitivity of the WT is 72.9%, of the DT 73.4% and the ST 23.3%, indicating that the distribution of the markers overlaps more in this state. The WT is mostly miss-classified as DT and vice versa, whereas ST is identified more as WT and DT than correctly as ST. To understand this further, the distributions for each class of the main spectral marker (i.e. fast




Fig. 6. Distribution of the power of the fast oscillations band for normal and stressed segments of each class: Wild Type (WT), Single transgenic PS2.30H (ST) and double transgenic B6.152H (DT).

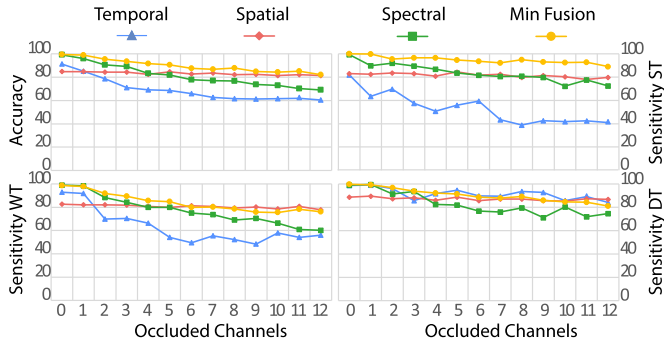


Fig. 7. Performance of the individual and ensemble models based on the number of replaced channels.

oscillations) are shown in Figure 6. For WT, the distribution has shifted to higher values while maintaining the same range. In the case of ST, the distribution has shifted to higher values, with the minimum also being higher. Lastly, the DT shows a higher power and a larger range. Overall, there is a higher overlap of the distributions, leading to incorrect classifications.

Finally, the ensemble model achieves a better performance than the temporal and spectral models but poorer performance than the spatial model due to the miss-classification of the ST.

D. Performance of Robustness to Artefacted Channels

The progress of the different performance metrics, such as accuracy and class sensitivities as an increasing number of channels is randomly replaced, are shown in Fig. 7. In regards to the accuracy, we see that while the spectral model starts with the lowest performance, it is the least affected one, as all the signals are averaged before feature extraction. On the other hand, the temporal model, which performed better, lost a significant amount of performance, especially in the first 3 removals. While the temporal model has a downward trend, it does show the most variability, which can be attributed to a bias to particular randomly occluded channels of the samples

that have been replaced with the mean. The spatial model also suffers from reduced performance, but lesser than the temporal model. Lastly, the min fusion model maintains a high accuracy throughout, with an accuracy over 90% even with 5 channels replaced.

Looking at the class-specific sensitivities, DT stands out as being the least affected of the three, meaning it has the most discriminative features. On the other hand, ST sensitivity decreases the most in the temporal model, whilst the WT of the spatial model. Overall, the spectral model maintains the performance due to what was previously mentioned. The min fusion model maintains high performance for the WT but is affected more than the spectral model for more than 6 replaced channels.

IV. CONCLUSION

This work was to develop a methodology that allows understanding AD via the discovery of biomarkers, for which a robust model was needed. A multimodal approach was used to label and structure the neuronal recordings, which subsequently had different features of each domain extracted from them as input for the classification models. Different combinations of ML models, diverse neural network architecture, and boosting methods were evaluated for each domain in order to find the best model for each one. Subsequently, a range of fusion techniques were explored to boost the overall performance. The final model achieved an accuracy of 99.4 % over the test set.

After analysing the insights provided by the individual models, it was discovered that the higher frequency bands, such as high gamma and fast oscillations, play a key role in the spectral domain. From the spatial models, it was validated that the hippocampus has an important role in AD detection. This biomarker can be translated to human patients via medical imaging techniques such as functional MRI [42], positron emission tomography (PET) scans [43] and computed tomography [44], that would allow comparing the mass, volume, or blood oxygenation level of the hippocampus to detect AD. Despite identifying the key areas of the brain for the detection of AD, a model trained with a pair of individual channels from the most relevant brain areas would not achieve the same performance as those trained with all channels, nor would it be robust to artefacts, thus we favour the multi-channel approach.

To evaluate the robustness of the models to segments where the rodents felt stressed, examples of these segments were shown to the models, where the classification performance dropped significantly. Results showed that the changes in the spectral properties of the signal due to stress were detrimental to the correct identification of AD, highlighting the importance of removing them during pre-processing. Subsequently, their performance was tested by randomly replacing channels in each example with up to half of the available ones. The ensemble's performance remained consistently satisfactory even with five channels replaced.

Overall, the proposed pipeline has proven to be both accurate and robust, as well as maintaining the ability to provide the feature insights needed for a deeper analysis.

REFERENCES

- [1] M. J. Prince et al., "World Alzheimer report 2015—the global impact of dementia: An analysis of prevalence, incidence, cost and trends," *Alzheimer's Disease Int.*, London, U.K., Tech. Rep. 2015, 2015, pp. 1–24.
- [2] W. Wong, "Economic burden of Alzheimer disease & managed care," *Amer. J. Managed Care*, vol. 26, no. 8, pp. S177–S183, 2020.
- [3] C. Cecere, C. Corrado, and R. Polikar, "Diagnostic utility of EEG based biomarkers for Alzheimer's disease," in *Proc. 40th Annu. Northeast Bioeng. Conf. (NEBEC)*, Apr. 2014, pp. 1–2.
- [4] J. Ruiz, M. Mahmud, M. Modasshir, M. S. Kaiser, and F. T. Alzheimer's Disease Neuroimaging Initiative, "3D DenseNet ensemble in 4-way classification of Alzheimer's disease," in *Brain Informatics*. Cham, Switzerland: Springer, 2020, pp. 85–96.
- [5] T. Chen et al., "Assessing significance of cognitive assessments for diagnosing Alzheimer's disease with fuzzy-rough feature selection," in *Proc. UKCI*, 2022, pp. 450–462.
- [6] S. Jahan et al., "Explainable AI-based Alzheimer's prediction and management using multimodal data," *Preprints*, pp. 1–16, 2022, doi: 10.20944/preprints202203.0214.v1.
- [7] N. Shaffi, F. Hajamohideen, M. Mahmud, A. Abdesselam, K. Subramanian, and A. A. Sariri, "Triplet-loss based Siamese convolutional neural network for 4-way classification of Alzheimer's disease," in *Brain Informatics*. Cham, Switzerland: Springer, 2022, pp. 277–287.
- [8] N. Shaffi, F. Hajamohideen, A. Abdesselam, M. Mahmud, and K. Subramanian, "Ensemble classifiers for a 4-way classification of Alzheimer's disease," in *Proc. AII*, 2023, pp. 219–230.
- [9] F. Hajamohideen et al., "Four-way classification of Alzheimer's disease using deep Siamese convolutional neural network with triplet-loss function," *Brain Inform.*, vol. 10, no. 1, pp. 1–13, 2023.
- [10] T. M. N. U. Akhund, M. J. N. Mahi, A. H. Tanvir, M. Mahmud, and M. S. Kaiser, "Adeptness: Alzheimer's disease patient management system using pervasive sensors—early prototype and preliminary results," in *Brain Informatics*. Cham, Switzerland: Springer, 2018, pp. 413–422.
- [11] S. Jesmin, M. S. Kaiser, and M. Mahmud, "Artificial and Internet of Healthcare Things based Alzheimer care during COVID 19," in *Brain Informatics*. Cham, Switzerland: Springer, 2020, pp. 263–274.
- [12] V. Vimbi, N. Shaffi, M. Mahmud, K. Subramanian, and F. Hajamohideen, "Application of explainable artificial intelligence in Alzheimer's disease classification: A systematic review," *Res. Square*, pp. 1–41, 2023, doi: 10.21203/rs.3.rs-2734771/v1.
- [13] M. Mahmud, M. S. Kaiser, A. Hussain, and S. Vassanelli, "Applications of deep learning and reinforcement learning to biological data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 6, pp. 2063–2079, Jun. 2018.
- [14] M. Mahmud, M. S. Kaiser, T. M. McGinnity, and A. Hussain, "Deep learning in mining biological data," *Cognit. Comput.*, vol. 13, no. 1, pp. 1–33, Jan. 2021.
- [15] M. Tanveer, B. Richhariya, R. U. Khan, A. H. Rashid, P. Khanna, M. Prasad, and T. C. Lin, "Machine learning techniques for the diagnosis of Alzheimer's disease: A review," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 16, no. 1s, pp. 1–35, Apr. 2020.
- [16] M. B. T. Noor, N. Z. Zenia, M. S. Kaiser, S. A. Mamun, and M. Mahmud, "Application of deep learning in detecting neurological disorders from magnetic resonance images: A survey on the detection of Alzheimer's disease, Parkinson's disease and schizophrenia," *Brain Informat.*, vol. 7, no. 1, pp. 1–21, Dec. 2020.
- [17] H. M. Ali, M. S. Kaiser, and M. Mahmud, "Application of convolutional neural network in segmenting brain regions from MRI data," in *Brain Informatics*. Cham, Switzerland: Springer, 2019, pp. 136–146.
- [18] M. Fabietti, M. Mahmud, and A. Lotfi, "Effectiveness of employing multimodal signals in removing artifacts from neuronal signals: An empirical analysis," in *Brain Informatics*. Cham, Switzerland: Springer, 2020, pp. 183–193.
- [19] J. D. R. Milan and J. Carmena, "Invasive or noninvasive: Understanding brain–machine interface technology [conversations in BME]," *IEEE Eng. Med. Biol. Mag.*, vol. 29, no. 1, pp. 16–22, Jan. 2010.
- [20] M. Fabietti et al., "SANTIA: A MATLAB-based open-source toolbox for artifact detection and removal from extracellular neuronal signals," *Brain Informat.*, vol. 8, no. 1, pp. 1–19, Dec. 2021.
- [21] S. J. Webster, A. D. Bachstetter, P. T. Nelson, F. A. Schmitt, and L. J. Van Eldik, "Using mice to model Alzheimer's dementia: An overview of the clinical disease and the preclinical behavioral changes in 10 mouse models," *Frontiers Genet.*, vol. 5, p. 88, Apr. 2014.
- [22] B. Laurijsens, F. Aujard, and A. Rahman, "Animal models of Alzheimer's disease and drug development," *Drug Discovery Today, Technol.*, vol. 10, no. 3, pp. e319–e327, 2013.
- [23] S. Beker, V. Kellner, G. Chechik, and E. A. Stern, "Learning to classify neural activity from a mouse model of Alzheimer's disease amyloidosis versus controls," *Alzheimer's Dementia, Diagnosis, Assessment Disease Monitor.*, vol. 2, no. 1, pp. 39–48, Jan. 2016.
- [24] A. Leparulo, M. Mahmud, E. Scremin, T. Pozzan, S. Vassanelli, and C. Fasolato, "Dampened slow oscillation connectivity anticipates amyloid deposition in the PS2APP mouse model of Alzheimer's disease," *Cells*, vol. 9, no. 1, p. 54, Dec. 2019.
- [25] M. Müller, "Dynamic time warping," in *Information Retrieval for Music and Motion*. Berlin, Germany: Springer, 2007, pp. 69–84.
- [26] J. R. Vergara and P. A. Estévez, "A review of feature selection methods based on mutual information," *Neural Comput. Appl.*, vol. 24, no. 1, pp. 175–186, Jan. 2014.
- [27] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "EEGNet: A compact convolutional neural network for EEG-based brain–computer interfaces," *J. Neural Eng.*, vol. 15, no. 5, Oct. 2018, Art. no. 056013.
- [28] M. Daianu et al., "Breakdown of brain connectivity between normal aging and Alzheimer's disease: A structural k -core network analysis," *Brain Connectivity*, vol. 3, no. 4, pp. 407–422, Aug. 2013.
- [29] A. Khan, A. Sohail, U. Zahoor, and A. S. Qureshi, "A survey of the recent architectures of deep convolutional neural networks," *Artif. Intell. Rev.*, vol. 53, no. 8, pp. 5455–5516, Dec. 2020.
- [30] R. A. Berk, "Classification and regression trees (CART)," in *Statistical Learning From a Regression Perspective*. New York, NY, USA: Springer, 2008, pp. 1–65.
- [31] T. G. Dietterich, "An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization," *Mach. Learn.*, vol. 40, no. 2, pp. 139–157, 2000.
- [32] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *J. Comput. Syst. Sci.*, vol. 55, no. 1, pp. 119–139, Aug. 1997.
- [33] N. K. Al-Qazzaz, Z. A. A. Alyasseri, K. H. Abdulkareem, N. S. Ali, M. N. Al-Mhiqani, and C. Guger, "EEG feature fusion for motor imagery: A new robust framework towards stroke patients rehabilitation," *Comput. Biol. Med.*, vol. 137, Oct. 2021, Art. no. 104799.
- [34] K. R. Babcock, J. S. Page, J. R. Fallon, and A. E. Webb, "Adult hippocampal neurogenesis in aging and Alzheimer's disease," *Stem Cell Rep.*, vol. 16, no. 4, pp. 681–693, Apr. 2021.
- [35] M. J. West, C. H. Kawas, L. J. Martin, and J. C. Troncoso, "The CA1 region of the human hippocampus is a hot spot in Alzheimer's disease," *Ann. New York Acad. Sci.*, vol. 908, no. 1, pp. 255–259, Jan. 2006.
- [36] Y. Xu, M. Zhao, Y. Han, and H. Zhang, "GABAergic inhibitory interneuron deficits in Alzheimer's disease: Implications for treatment," *Frontiers Neurosci.*, vol. 14, p. 660, Jun. 2020.
- [37] K. Kalemaki, X. Konstantoudaki, S. Tivodar, K. Sidiropoulou, and D. Karageorgos, "Mice with decreased number of interneurons exhibit aberrant spontaneous and oscillatory activity in the cortex," *Frontiers Neural Circuits*, vol. 12, p. 96, Oct. 2018.
- [38] M. F. Carr, M. P. Karlsson, and L. M. Frank, "Transient slow gamma synchrony underlies hippocampal memory replay," *Neuron*, vol. 75, no. 4, pp. 700–713, Aug. 2012.
- [39] L. E. Arroyo-García et al., "Impaired spike-gamma coupling of area CA3 fast-spiking interneurons as the earliest functional impairment in the App^{NL-GF} mouse model of Alzheimer's disease," *Mol. Psychiatry*, vol. 26, no. 10, pp. 5557–5567, Oct. 2021.
- [40] D. Swietlik, J. Bialowkas, J. Morys, I. Klejbor, and A. Kusiak, "Effects of inducing gamma oscillations in hippocampal subregions DG, CA3, and CA1 on the potential alleviation of Alzheimer's disease-related pathology: Computer modeling and simulations," *Entropy*, vol. 21, no. 6, p. 587, Jun. 2019.
- [41] J. A. van Deursen, E. F. P. M. Vuurman, F. R. J. Verhey, V. H. J. M. van Kranen-Mastenbroek, and W. J. Riedel, "Increased EEG gamma band activity in Alzheimer's disease and mild cognitive impairment," *J. Neural Transmiss.*, vol. 115, no. 9, pp. 1301–1311, Sep. 2008.
- [42] L. Wang et al., "Changes in hippocampal connectivity in the early stages of Alzheimer's disease: Evidence from resting state fMRI," *NeuroImage*, vol. 31, no. 2, pp. 496–504, Jun. 2006.
- [43] S. Yamaguchi et al., "Decreased cortical glucose metabolism correlates with hippocampal atrophy in Alzheimer's disease as shown by MRI and PET," *J. Neurol., Neurosurgery Psychiatry*, vol. 62, no. 6, pp. 596–600, Jun. 1997.
- [44] E. Porter et al., "Hippocampus segmentation on noncontrast CT using deep learning," *Med. Phys.*, vol. 47, no. 7, pp. 2950–2961, Jul. 2020.