Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Learnable Multiresolution Graph Representations for EEG-Based Alzheimer's Disease and Frontotemporal Dementia Classification

Adi Alhudhaif

a.alhudhaif@psau.edu.sa

Prince Sattam bin Abdulaziz University

Ahmad O. Aseeri

Prince Sattam bin Abdulaziz University

Research Article

Keywords: EEG, Alzheimer's disease, Frontotemporal dementia, Learnable graph signal processing, Graph wavelet transform, Subject-independent validation

Posted Date: November 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11772642/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Electroencephalography (EEG) offers a non-invasive means of detecting neurophysiological changes associated with Alzheimer's disease (AD) and frontotemporal dementia (FTD), but the evaluation protocol determines whether reported performance reflects discrimination among segments of previously seen participants or generalization to unseen individuals. This study proposes a learnable multiresolution graph framework for EEG classification and evaluates it under segment-level and subject-independent partitioning. The framework jointly optimizes three branches: graph Fourier processing with learnable Chebyshev spectral filters, graph wavelet processing with learnable scales, and wavelet-like temporal convolution. A weighted electrode graph is constructed from absolute Pearson correlations among 19 EEG channels, with graph spectral processing defined through the normalized Laplacian. Branch features are concatenated, standardized and classified by a multilayer perceptron. Experiments use the OpenNeuro ds004504 dataset, comprising 36 participants with AD, 23 with FTD and 29 healthy controls (HC). Under segment-level partitioning, accuracies were 98.96% for AD vs. HC, 98.90% for AD vs. FTD, 99.13% for FTD vs. HC and 98.48% for three-class classification. Under subject-independent five-fold cross-validation, in which all segments of a participant are assigned to a single fold, AD vs. HC classification achieved an accuracy of 84.65 ± 1.06% and an F1-score of 80.33 ± 1.06%, 14.31 percentage points below the segment-level accuracy. The framework learns complementary graph spectral and temporal representations, and the gap between the two protocols shows that participant-level separation is essential when estimating performance on unseen individuals.

1 Introduction

Neurodegenerative diseases, including Alzheimer’s disease (AD) and frontotemporal dementia (FTD), afflict more than 55 million individuals globally, with an estimated 10 million new cases each year [1, 2]. AD is responsible for as many as 70% of these individuals, making it the most prevalent etiology of dementia [3]. AD usually begins with memory problems, whereas FTD begins with changes in behavior and speech. EEG, an inexpensive and non-invasive method of monitoring brain activity, is a promising way to detect the early changes linked to these disorders [4, 5]. AD is characterized by a generalized slowing of cortical rhythms, with increased delta and theta activity and reduced posterior alpha and beta activity, reflecting diffuse thalamo-cortical and cortico-cortical dysfunction. FTD typically preserves the posterior alpha rhythm but shows reduced beta power, irregular frontal patterns and altered frontotemporal connectivity, whereas healthy controls (HC) show distinct posterior alpha peaks with low delta and theta power.

Machine learning (ML) and deep learning (DL) have been applied extensively to EEG-based dementia classification. Lightweight convolutional neural networks (CNNs) include LEADNet [6], which reached 99.24% test accuracy on 12 AD patients and 11 controls, STEADYNet [7], which reported 99.29% for AD, 99.65% for MCI and 92.25% for FTD, the spectrogram-based CNN of Stefanou et al. [8] (98.38%) and LCADNet [9]. Feature-based pipelines combined time, frequency and complexity descriptors with LightGBM [10], optimal channel selection with the tunable Q-wavelet transform [11] and optimal wavelet

filter banks [12]. Hybrid architectures fused CNNs with Vision Transformers [13] (80.23% for AD/FTD/HC), CNN and Transformer encoders in DICE-Net [14] (83.28% for AD vs. HC under leave-one-subject-out validation), ConvNeXt-style networks on wavelet spectrograms in EEGConvNeXt [15] (95.70% for three classes) and multimodal multi-stage time-frequency representations [16]. Spectral studies showed that oscillatory [17] and aperiodic [18] EEG components distinguish AD, FTD and HC.

Graph and wavelet representations have also proved effective for other EEG and neuroimaging problems: graph signal processing and graph learning for schizophrenia [19], attention-based wavelet CNNs for epilepsy [20] and graph features for AD detection [21]. In most of these works, however, the graph or wavelet transform is fixed in advance, and most results on the AD/FTD dataset used here were obtained with segment-level splitting, in which windows of the same participant appear in both training and test sets.

This paper presents an EEG classification model for AD, FTD and HC in which three learnable modules, the graph Fourier transform (GFT), the graph wavelet transform (GWT) and a wavelet-like temporal convolution (referred to as DWT), feed a multilayer perceptron (MLP). The main contributions are:

1. An end-to-end trainable multiresolution framework in which the GFT, the GWT and a wavelet-like temporal convolution are optimized jointly with the classifier; the spectral filters and the wavelet scales are learned from the data rather than fixed in advance.
2. Graph-spatial, graph-spectral and temporal representations are combined inside a single differentiable architecture; an ablation study isolates what each learnable component contributes.
3. The same model, preprocessing and data are evaluated under two explicitly defined protocols, segment-level and subject-independent, separating the contribution of the protocol from that of the model.
4. The optimism introduced by segment-level splitting on the AD vs. HC task is quantified at 14.31 accuracy points and 18.51 F1 points, and the literature comparison for this dataset is organized by split unit.

2 Materials and methods

2.1 Overview of the proposed framework

Figure 1 illustrates the proposed framework. The raw EEG is band-pass filtered and reshaped, and a brain graph is constructed from the absolute Pearson correlation coefficients of each pair of channels, which gives a fully connected weighted adjacency matrix with zeros on the diagonal. No hard threshold is applied: thresholding would add a hyperparameter, and weak but non-zero channel relationships may carry the diffuse topological changes seen in neurodegenerative disease. Instead, the learnable spectral filters in the GFT and GWT modules suppress uninformative graph frequencies, so sparsity is learned in the graph spectrum rather than imposed on the graph. Pearson correlation remains sensitive to volume conduction, so the resulting matrix is better described as a data-driven support for the learnable GFT and

GWT layers than as a model of anatomical connectivity (Section 4). The network (EEGNet) integrates a GFT block for graph-spectral features, a temporal convolutional block in the spirit of the DWT and a multiresolution GWT over the electrode graph. Their outputs are concatenated, normalized and passed to an MLP with one fully connected layer of 128 units, ReLU activation and dropout, followed by an output layer. Class weights computed on the training partition counteract label imbalance, and the train-test split follows one of the protocols of Section 2.7.

2.2 Learnable graph Fourier transform

Each node of the graph corresponds to an EEG channel and edges are weighted by correlation [21]. The normalized graph Laplacian is

$$L = I - D^{-1/2}\, W\, D^{-1/2},$$

1

where W is the adjacency matrix, D the diagonal degree matrix and I the identity matrix. Its eigendecomposition $L = U\Lambda\, U^T$ gives the graph Fourier basis U and the graph frequencies Λ, and a graph signal x is projected into the spectral domain as $\widehat{x} = U^T x$. In the learnable variant the spectral filters are parameterized and learned from data:

$$T\,(x) = U\, g(\Lambda\,;\theta\,)\, U^T x,$$

2

where $g(\Lambda\,;\theta\,)$ is a spectral kernel applied element-wise to the eigenvalues, with parameters θ optimized during training. The filtering is implemented with Chebyshev polynomial approximations, which avoid repeated eigendecompositions and support localized operations. The filter response is not constrained in advance, so the module selects the graph frequencies that are most indicative of AD, FTD and HC.

2.3 Learnable temporal wavelet transform

The learnable wavelet transform learns convolutional filters that simulate wavelet-like decompositions [22]. Let $x \in \mathbb{R}^{N\times\, T\times\, C}$ denote the input, where N is the number of samples, T the number of time samples and C the number of channels. Each channel x_c is processed by F learned filters $h_f^{(1)}$ of size k, followed by a ReLU $\sigma\,(\cdot\,)$ and a second convolution layer:

$$z_f^{(1)}(t,c) = \sigma\,\left(\sum_{\tau\,=-\lfloor\,k/2\rfloor}^{\lfloor\,k/2\rfloor} h_f^{(1)}\,(\tau\,)\; x_c(t-\tau\,)\right),$$

3

$$z_f^{(2)}(t,c) = \sigma\,\left(\sum_{\tau\,=-\lfloor\,k/2\rfloor}^{\lfloor\,k/2\rfloor} h_f^{(2)}\,(\tau\,)\; z_f^{(1)}(t-\tau\,,c)\right).$$

4

Composing local filters into broader receptive fields captures fine and coarse temporal patterns. Global average pooling over time and averaging over channels give a fixed-size embedding:

$$w_f = \frac{1}{CT} \sum_{c=1}^{C} \sum_{t=1}^{T} z_f^{(2)}(t, c),$$

$$w = [w_1, \ldots, w_F].$$

5

2.4 Learnable graph wavelet transform

The GWT performs localized multiscale filtering in the spectral domain of the graph [22, 23]. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with $C = |\mathcal{V}|$ nodes carrying a signal $x \in \mathbb{R}^C$, with Laplacian $L \in \mathbb{R}^{C \times C}$ decomposed as in Section 2.2. For a learnable scale $s \in \mathbb{R}_+$ the wavelet kernel is

$$g_s(\lambda) = \left(1 - (s\lambda)^2\right) \exp\left(-0.5\,(s\lambda)^2\right),$$

6

which emphasizes intermediate frequencies while suppressing low and high ones [23]. For each of L learnable scales $\{s_1, \ldots, s_L\}$ the kernel is evaluated at the eigenvalues, the spectral coefficients are modulated and transformed back to the nodes, and the results are concatenated:

$$x^{(s_l)} = U\left(\widehat{x} \odot g_{s_l}\right),$$

7

$$x_{\mathrm{GWT}} = \left[x^{(s_1)} \| \cdots \| x^{(s_L)}\right] \in \mathbb{R}^{LC}.$$

8

Unlike the Mexican-hat or Hammond–Vandergheynst kernels, the scales are trained jointly with the GFT, DWT and MLP modules, so the scale-frequency responses that separate AD, FTD and HC are learned from the task.

2.5 Statistical feature extraction

Interpretable statistical descriptors are computed per channel. For the univariate series $x_c(t)$ of length T, the mean μ_c, variance σ_c^2, skewness γ_c, kurtosis k_c, normalized energy E_c and Shannon entropy H_c (estimated from M histogram bins with probabilities p_i) are

$$\mu_c = \frac{1}{T}\sum_{t=1}^{T} x_c(t),$$

$$\sigma_c^2 = \frac{1}{T}\sum_{t=1}^{T} (x_c(t) - \mu_c)^2,$$

$$\gamma_c = \frac{1}{T}\sum_{t=1}^{T} \left(\frac{x_c(t) - \mu_c}{\sigma_c}\right)^3,$$

$$k_c = \frac{1}{T}\sum_{t=1}^{T} \left(\frac{x_c(t) - \mu_c}{\sigma_c}\right)^4,$$

$$E_c = \frac{1}{T}\sum_{t=1}^{T} x_c(t)^2,$$

$$H_c = -\sum_{i=1}^{M} p_i \log p_i.$$

9

2.6 Dataset

The EEG dataset was recorded at the Second Department of Neurology, AHEPA University Hospital, Thessaloniki, Greece [24], and is publicly available as OpenNeuro ds004504. It contains 88 subjects: 36 patients with AD, 23 with FTD and 29 HC, recorded in a resting, eyes-closed state with a Nihon Kohden EEG-2100 system using 19 scalp electrodes (international 10–20 system, mastoid reference A1–A2). The signals were sampled at 500 Hz with electrode impedance below 5 k Ω , and sessions lasted about 13 minutes on average. The released data were re-referenced, band-pass filtered (0.5–45 Hz) and cleaned with Artifact Subspace Reconstruction and ICA with ICLabel in EEGLAB, and are stored in BIDS format.

2.7 Evaluation protocols

Both protocols use the identical model, preprocessing, hyperparameters and features; only the unit of the train-test split differs. Under the *segment-level protocol*, each recording is divided into non-overlapping windows, the windows of all participants are pooled and split at random into 80% training and 20% test data, so a participant contributes windows to both partitions. This is the protocol used in most published work on this dataset and is reported for comparability. Under the *subject-independent protocol*, the participants are partitioned into five class-stratified folds, all windows of a participant belong to one fold, and the mean and standard deviation over the five folds are reported.

The distinction matters because EEG carries strong participant-specific information unrelated to diagnosis: skull and scalp geometry, electrode impedance, individual alpha frequency, medication and session amplitude scaling. Under the segment-level protocol a classifier can score highly by recognizing which participant produced a window. A segment-level score is therefore an upper bound on what the

representation can separate, whereas the subject-independent score is the quantity relevant to clinical use and is treated here as the main result.

3 Results

3.1 Implementation details

The method was implemented in PyTorch and run on an NVIDIA GeForce RTX 3090 GPU (24 GB). Signals were band-pass filtered between 1 and 50 Hz, segmented into non-overlapping windows of 5000 samples (10 s at 500 Hz) unless stated otherwise and z-score normalized. All parameters of the GFT, DWT and GWT modules were learnable. The MLP uses 128 hidden units, ReLU and dropout of 0.5. Training used AdamW with a learning rate of 0.001, a batch size of 32, a StepLR scheduler halving the learning rate every 40 epochs, a weighted cross-entropy loss with weights computed on the training set, and 50 epochs. Training always ran for the fixed 50 epochs with the fixed learning-rate schedule, and all reported values are those of the final epoch. The test partition was evaluated after each epoch only to plot the learning curves; it was not used for model selection, early stopping or hyperparameter tuning. Performance is reported as accuracy, precision, recall, specificity and F1-score (%). Four tasks were studied: the three-class problem and the pairwise problems AD vs. HC, AD vs. FTD and FTD vs. HC. Sections 3.2 and 3.3 report segment-level results, and Section 3.4 reports the subject-independent result and compares the two protocols.

3.2 Classification under the segment-level protocol

For all four tasks, training and test accuracy rose sharply during the early epochs, exceeded 95% within the first 10 epochs and stabilized around 98–100%, with the training and test curves staying close. The confusion matrices are shown in Fig. 2. In the three-class task, 563 of the 574 AD test segments, 486 of the 495 HC segments and 315 of the 316 FTD segments are classified correctly. For AD vs. HC, 577 AD and 469 HC segments are classified correctly, with 7 AD segments labeled HC and 4 HC segments labeled AD. For AD vs. FTD, 593 of the 597 AD segments and 303 of the 307 FTD segments are correct. For FTD vs. HC, 801 of 808 test segments are correct: 5 HC segments are labeled FTD and 2 FTD segments are labeled HC, although the EEG differences between FTD and HC are subtler than those between AD and HC.

Table 1
Performance under the segment-level protocol (MLP classifier, 5000-sample windows unless stated otherwise)

Task	Accuracy (%)	Precision (%)	Recall (%)	Specificity (%)	F1-score (%)
Three-class, hold-out	98.48	98.20	98.65	98.43	98.41
Three-class, 10000-sample windows	96.91	96.77	96.73	96.75	96.75
Three-class, 10-fold CV	97.83 ± 0.50	97.74 ± 0.47	97.86 ± 0.54	97.80 ± 0.49	97.79 ± 0.50
AD vs. HC (hold-out)	98.96	98.53	99.15	98.80	98.84
AD vs. FTD (hold-out)	98.90	98.06	98.70	99.00	98.38
FTD vs. HC (hold-out)	99.13	98.37	99.34	99.01	98.85

Table 1 summarizes the metrics. The three-class task reaches 98.48% accuracy and a 98.41% F1-score with 5000-sample windows. With 10000-sample windows (3397 input samples) accuracy drops to 96.91%, so the shorter window is preferable; 10-fold cross-validation on segments gives 97.83 ± 0.50%. The binary tasks reach 98.96% (AD vs. HC), 98.90% (AD vs. FTD) and 99.13% (FTD vs. HC) accuracy, with specificities between 98.80% and 99.01%. The areas under the ROC curve are 0.997 (AD vs. FTD), 0.999 (AD vs. HC), 1.000 (FTD vs. HC) and 0.999 (three-class, micro-averaged one-vs-rest), and the corresponding average precision values are 0.997, 1.000, 0.999 and 0.997.

3.3 Ablation study

To assess the learnable components, the GFT and GWT modules were replaced by fixed counterparts (Table 2, three-class task, segment-level protocol). With both transforms fixed, the framework reaches 97.26% accuracy and a 97.21% F1-score, reflecting the information already present in the statistical descriptors and the temporal branch. A learnable GFT with a fixed GWT raises accuracy to 98.27% (+ 1.01 points), a learnable GWT alone gives 97.83%, and both learnable give 98.48%. The task-adaptive graph spectral filter accounts for most of the improvement, and the two learnable modules are complementary rather than redundant.

Table 2
Ablation study on the three-class task (segment-level protocol); all values in %

GFT	GWT	Accuracy	Precision	Recall	F1
Basic	Basic	97.26	97.17	97.25	97.21
Learnable	Basic	98.27	98.12	98.22	98.17
Basic	Learnable	97.83	97.75	97.91	97.83
Learnable	Learnable	98.48	98.20	98.41	98.41

3.4 Subject-independent evaluation and effect of the protocol

To estimate how the framework behaves on participants it has never seen, the AD vs. HC task was repeated under the subject-independent protocol (Table 3). The fold accuracies are 85.40%, 85.12%, 84.74%, 85.40% and 82.61% (F1-scores 80.37%, 80.00%, 80.64%, 81.92% and 78.70%), giving a mean accuracy of 84.65% with a standard deviation of 1.06%; mean precision, recall and F1-score are 79.96%, 80.71% and 80.33%, and mean specificity is 87.93%. The fold-to-fold spread is small, but each test fold contains only about 13 participants.

Table 3

Effect of the evaluation protocol on AD vs. HC classification (subject-independent values are mean ± std over five folds). The model, preprocessing, hyperparameters and data are identical; only the unit of the train-test split differs

Subject-independent 5-fold (mean ± std)	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**Specificity (%)**	**F1-score (%)**
	84.65 ± 1.06	**79.96 ± 1.16**	**80.71 ± 1.33**	**87.93 ± 1.74**	**80.33 ± 1.06**
Segment-level hold-out (Table 1)	98.96	98.53	99.15	98.80	98.84
Difference	−14.31	−18.57	−18.44	−10.87	−18.51

The bottom rows of Table 3 place the two protocols side by side. Moving from segment-level to subject-independent splitting costs 14.31 accuracy points, 18.57 precision points, 18.44 recall points and 18.51 F1 points. Specificity falls least, by 10.87 points, consistent with healthy controls forming a more homogeneous group than AD patients: the model still recognizes a typical control recording on an unseen participant, while the AD group varies more between individuals. The size of this gap is the main empirical result of the study: near-perfect accuracy on this dataset is largely a property of the split unit

rather than of the representation, and an accuracy above 95% cannot be interpreted as performance on a new patient unless the split unit is stated.

3.5 Explainability with Shapley analysis

A SHAP (Shapley Additive Explanations) analysis [25] was carried out by averaging the absolute SHAP values over the test set (Fig. 3; 0: AD, 1: HC, 2: FTD). In the 194-dimensional feature vector the first 16 entries come from the GFT, the next 64 from the DWT branch and the remaining 114 from the GWT branch. The largest contributions come from Var6, Var28, Var56, Var25 and Var30: four of these fall in the DWT block and one in the GFT block, so the temporal multiscale representation and the low-order graph-spectral components carry most of the discriminative information, while the GWT contributes more diffusely. This ordering agrees with the ablation, where a learnable GFT gave a larger gain (1.01 points) than a learnable GWT (0.57 points).

In the class-wise SHAP summary plots, for AD, Var6, Var3 (GFT) and Var25, Var56 (DWT) have the highest impact, so both graph-spectral and time-frequency representations contribute. For HC, Var56, Var28, Var59 (DWT) and Var6 (GFT) dominate with smaller magnitudes, and several features act in both directions, suggesting non-linear interactions. FTD has the most spread-out distribution, led by Var6 (GFT) and Var73, Var57, Var74 (DWT). Overall, AD predictions draw more on low-frequency and connectivity-related features, consistent with EEG slowing and network disruption in AD, whereas FTD predictions draw more on mid-frequency and localized patterns, consistent with focal frontotemporal pathology; attributions for HC are more evenly distributed.

Projecting the average EEG signal of a test window onto the Laplacian eigenbasis and evaluating the learned filters at the graph frequencies further showed that the learned filters vary smoothly and concentrate where the graph Fourier coefficients of the signal are concentrated, whereas a fixed graph Fourier kernel produces a sharply decaying response confined to low graph frequencies.

4 Discussion

Table 4
Comparison with previous studies on the same dataset. Figures obtained with different split units are not directly comparable

Authors	Year	Method	Results (%)
Acharya et al. [15]	2025	EEGConvNeXt	AD–HC: 96.32, AD–FTD: 98.21, FTD–HC: 98.67, AD–FTD–HC: 95.70
Amer et al. [26]	2024	GoogleNet, AlexNet, SqueezeNet	AD–HC: 95.91, 93.03, 91.72
Zheng et al. [27]	2024	SVM	AD–HC: 87.69, FTD–HC: 82.69, AD–FTD: 72.88, HC–(AD&FTD): 86.36
Kachare et al. [7]	2024	STEADYNet	AD–HC: 88.00, AD–FTD: 89.40, FTD–HC: 92.25
Zikereya et al. [17]	2024	SVM	AD–HC: 83.00, FTD–HC: 70.00, AD–FTD: 60.00
Wang et al. [18]	2024	SVM	AD–FTD: 73.00
Miltiadous et al. [10]	2023	LightGBM with complexity features	AD–HC: 83.06 (F1), FTD–HC: 80.67 (F1), Dementia–HC: 92.27 (F1)
Miltiadous et al. [14]	2023	DICE-Net (CNN + Transformer)	AD–HC: 83.28, FTD–HC: 74.43
This study (segment-level)	2026	Learnable multiresolution graph signal processing	AD–HC: 98.96, AD–FTD: 98.90, FTD–HC: 99.13, AD–FTD–HC: 98.48
This study (subject-independent)	2026	Learnable multiresolution graph signal processing	AD–HC: 84.65 ± 1.06

Table 4 compares this work with earlier studies on the same dataset. Acharya et al. [15] report 95.70% for the three-class task with EEGConvNeXt, Amer et al. [26] up to 95.91% for AD-HC with pretrained CNNs, and Kachare et al. [7] 88.00–92.25% on the pairwise tasks. SVM and ensemble methods report lower figures [17, 18, 27], with AD-FTD accuracy as low as 60.00%, and Miltiadous et al. [10, 14] obtained roughly 71–92% depending on the classifier and features. Our segment-level accuracies exceed every figure in the table, but that comparison is not the informative one, because several cited results were obtained under subject-independent protocols. Our subject-independent AD vs. HC accuracy of 84.65 ± 1.06% is slightly above the 83.28% of DICE-Net [14] under leave-one-subject-out validation, the highest subject-independent accuracy on this dataset that we are aware of. The improvement is real but modest.

The gap between the two protocols is the more consequential finding. Changing only the split unit costs 14.31 accuracy points and 18.51 F1 points on AD vs. HC (Table 3). Unless the split unit is stated, an

accuracy above 95% says how well a model separates windows, not patients; we encourage future work on this dataset to report both numbers.

The main strengths of the method are that (i) the transforms are learned rather than prescribed: the GFT filter coefficients, GWT scales and temporal kernels are optimized together with the classifier, which raises three-class accuracy from 97.26% with fixed transforms to 98.48% (Table 2); (ii) the same architecture handles the binary and three-class problems under both protocols without task-specific redesign; and (iii) the model is compact, with a single hidden layer of 128 units acting on a 194-dimensional feature vector of 10 s windows, which is modest compared with transformer and pretrained-CNN pipelines.

The study also has limitations. (i) Subject-independent validation was carried out only for AD vs. HC; the other results come from segment-level splitting and should be read as upper bounds. (ii) A single dataset of 88 participants from one site and one acquisition system was used, each subject-independent test fold holds about 13 participants, and no external cohort was available. (iii) The folds were stratified by class but not matched for age, sex or MMSE score, so part of the subject-independent performance may reflect these variables. (iv) Absolute Pearson correlation is sensitive to volume conduction; phase-lag index, phase-locking value or mutual information would be less affected but were not compared. (v) Band-pass filtering and window length remain manually chosen; doubling the window from 5000 to 10000 samples lowered three-class accuracy from 98.48% to 96.91%. (vi) The model is fully supervised and requires labeled recordings.

5 Conclusions and future work

This study developed a learnable multiresolution graph signal processing framework in which the GFT, the GWT and a wavelet-like temporal convolution are trained jointly with the classifier. Under segment-level splitting it reaches 98.48% accuracy on the three-class task and 98.90–99.13% on the binary tasks; under subject-independent splitting, AD vs. HC accuracy is $84.65 \pm 1.06\%$. The learnable representation gives a 1.22-point gain over fixed transforms and a subject-independent accuracy slightly above the best previously published result on this dataset, while the 14.31-point gap between protocols shows that accuracies above 98% describe how well windows, not patients, are separated. Future work will extend subject-independent evaluation to the AD vs. FTD, FTD vs. HC and three-class tasks, validate the framework on an independent cohort, replace Pearson correlation with a connectivity measure robust to volume conduction, and match folds on age and MMSE.

Declarations

Competing interests:

The authors declare no competing interests.

Ethics approval and consent to participate:

This study was conducted using a publicly available dataset. Therefore, ethical approval from an Institutional Review Board was not required, and patient consent was not required because the study was based on secondary analysis of a publicly available open-access dataset.

Consent for publication:

Not applicable.

Funding:

This research was funded by Prince Sattam bin Abdulaziz University through project number PSAU/2026/01/32652. The funders had no role in the design of the study; in the collection, analyses or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Author Contribution

All authors have equal contributions in this work.

Acknowledgements

The authors extend their appreciation to Prince Sattam bin Abdulaziz University for funding this research work through project number PSAU/2026/01/32652.

Data availability:

The dataset analyzed in this study is publicly available at https://openneuro.org/datasets/ds004504.

References

1. Wimo, A., Prince, M.J.: World Alzheimer Report 2010: The Global Economic Impact of Dementia. Technical report, Alzheimer's Disease International, London (2010)
2. Nooreen, Z., Sumayya, F., Wal, P., Rai, A.K., Tandon, S.: Frontotemporal dementia, current diagnosis and treatment emphasis on natural products: A review. Curr. Aging Sci. (2025)
3. Prince, M., Wimo, A., Guerchet, M., Ali, G.C., Wu, Y.T., Prina, M.: World Alzheimer Report 2015. The Global Impact of Dementia: An Analysis of Prevalence, Incidence, Cost and Trends. Technical report, Alzheimer's Disease International, London (2015)
4. Sawangjai, P., Hompoonsup, S., Leelaarporn, P., Kongwudhikunakorn, S., Wilaiprasitporn, T.: Consumer grade EEG measuring sensors as research tools: A review. IEEE Sens. J. **20**(8), 3996–4024 (2020). https://doi.org/10.1109/JSEN.2019.2962874

5. Ahmed, M.R., Zhang, Y., Feng, Z., Lo, B., Inan, O.T., Liao, H.: Neuroimaging and machine learning for dementia diagnosis: Recent advancements and future prospects. IEEE Rev. Biomed. Eng. **12**, 19–33 (2018)
6. Puri, D.V., Kachare, P.H., Sangle, S.B., Kirner, R., Jabbari, A., Al-Shourbaji, I., Abdalraheem, M., Alameen, A.: LEADNet: Detection of Alzheimer's disease using spatiotemporal EEG analysis and low-complexity CNN. IEEE Access. **12**, 113888–113897 (2024). https://doi.org/10.1109/ACCESS.2024.3435768
7. Kachare, P.H., Sangle, S.B., Puri, D.V., Khubrani, M.M., Al-Shourbaji, I.: STEADYNet: Spatiotemporal EEG analysis for dementia detection using convolutional neural network. Cogn. Neurodyn. **18**(5), 3195–3208 (2024)
8. Stefanou, K., Tzimourta, K.D., Bellos, C., Stergios, G., Markoglou, K., Gionanidis, E., Miltiadous, A., et al.: A novel CNN-based framework for Alzheimer's disease detection using EEG spectrogram representations. J. Personalized Med. **15**(1), 27 (2025)
9. Kachare, P., Puri, D., Sangle, S.B., Al-Shourbaji, I., Jabbari, A., Kirner, R., Alameen, A., Migdady, H., Abualigah, L.: LCADNet: A novel light CNN architecture for EEG-based Alzheimer disease detection. Phys. Eng. Sci. Med. **47**(3), 1037–1050 (2024). https://doi.org/10.1007/s13246-024-01425-w
10. Miltiadous, A., Tzimourta, K.D., Aspiotis, V., Afrantou, T., Tsipouras, M.G., Giannakeas, N., Glavas, E., Tzallas, A.T.: Enhanced Alzheimer's disease and frontotemporal dementia EEG detection: Combining LightGBM gradient boosting with complexity features. In: Proceedings of the 2023 IEEE 36th International Symposium on Computer-Based Medical Systems (CBMS), pp. 876–881 (2023). https://doi.org/10.1109/CBMS58004.2023.00335
11. Puri, D., Nalbalwar, S., Nandgaonkar, A., Wagh, A.: Alzheimer's disease detection from optimal electroencephalogram channels and tunable Q-wavelet transform. Indonesian J. Electr. Eng. Comput. Sci. **25**(3), 1420–1428 (2022)
12. Puri, D.V., Gawande, J.P., Rajput, J.L., Nalbalwar, S.L.: A novel optimal wavelet filter banks for automated diagnosis of Alzheimer's disease and mild cognitive impairment using electroencephalogram signals. Decis. Analytics J. **9**, 100336 (2023). https://doi.org/10.1016/j.dajour.2023.100336
13. Chen, Y., Wang, H., Zhang, D., Zhang, L., Tao, L.: Multi-feature fusion learning for Alzheimer's disease prediction using EEG signals in resting state. Front. NeuroSci. **17**, 1272834 (2023). https://doi.org/10.3389/fnins.2023.1272834
14. Miltiadous, A., Gionanidis, E., Tzimourta, K.D., Giannakeas, N., Tzallas, A.T.: DICE-net: A novel convolution-transformer architecture for Alzheimer detection in EEG signals. IEEE Access. **11**, 71840–71858 (2023)
15. Acharya, M., Deo, R.C., Barua, P.D., Devi, A., Tao, X.: EEGConvNeXt: A novel convolutional neural network model for automated detection of Alzheimer's disease and frontotemporal dementia using EEG signals. Comput. Methods Programs Biomed. **262**, 108652 (2025). https://doi.org/10.1016/j.cmpb.2025.108652

16. Vo, T., Ibrahim, A.K., Zhuang, H.: A multimodal multi-stage deep learning model for the diagnosis of Alzheimer's disease using EEG measurements. Neurol. Int. **17**(6), 91 (2025)
17. Zikereya, T., Lin, Y., Zhang, Z., Taguas, I., Shi, K., Han, C.: Different oscillatory mechanisms of dementia-related diseases with cognitive impairment in closed-eye state. NeuroImage. **304**, 120945 (2024)
18. Wang, Z., Liu, A., Yu, J., Wang, P., Bi, Y., Xue, S., Zhang, W., et al.: The effect of aperiodic components in distinguishing Alzheimer's disease from frontotemporal dementia. GeroScience. **46**(1), 751–768 (2024)
19. Pain, S., Sarma, M., Samanta, D.: Graph signal processing and graph learning approaches to schizophrenia pattern identification in brain electroencephalogram. Biomed. Signal Process. Control. **100**, 106954 (2025)
20. Xin, Q., Hu, S., Liu, S., Zhao, L., Zhang, Y.-D.: An attention-based wavelet convolution neural network for epilepsy EEG classification. IEEE Trans. Neural Syst. Rehabil. Eng. **30**, 957–966 (2022). https://doi.org/10.1109/TNSRE.2022.3166181
21. Sharma, R., Meena, H.K.: Graph based novel features for detection of Alzheimer's disease using EEG signals. Biomed. Signal Process. Control. **103**, 107380 (2025). https://doi.org/10.1016/j.bspc.2024.107380
22. Hammond, D.K., Vandergheynst, P., Gribonval, R.: The spectral graph wavelet transform: Fundamental theory and fast computation. In: Stanković, L., Sejdić, E. (eds.) Vertex-Frequency Analysis of Graph Signals. Signals and Communication Technology, pp. 141–175. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-03574-7_3
23. Hammond, D.K., Vandergheynst, P., Gribonval, R.: Wavelets on graphs via spectral graph theory. Appl. Comput. Harmon. Anal. **30**(2), 129–150 (2011). https://doi.org/10.1016/j.acha.2010.04.005
24. Miltiadous, A., Tzimourta, K.D., Afrantou, T., Ioannidis, P., Grigoriadis, N., Tsalikakis, D.G., Tzallas, A.T., et al.: A dataset of scalp EEG recordings of Alzheimer's disease, frontotemporal dementia and healthy subjects from routine EEG. Data. **8**(6), 95 (2023)
25. Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems, vol. 30 (2017)
26. Amer, N.S., Belhaouari, S.B.: Exploring new horizons in neuroscience disease detection through innovative visual signal analysis. Sci. Rep. **14**(1), 4217 (2024)
27. Zheng, H., Xiong, X., Zhang, X.: Multi-threshold recurrence rate plot: A novel methodology for EEG analysis in Alzheimer's disease and frontotemporal dementia. Brain Sci. **14**(6), 565 (2024)

Figures

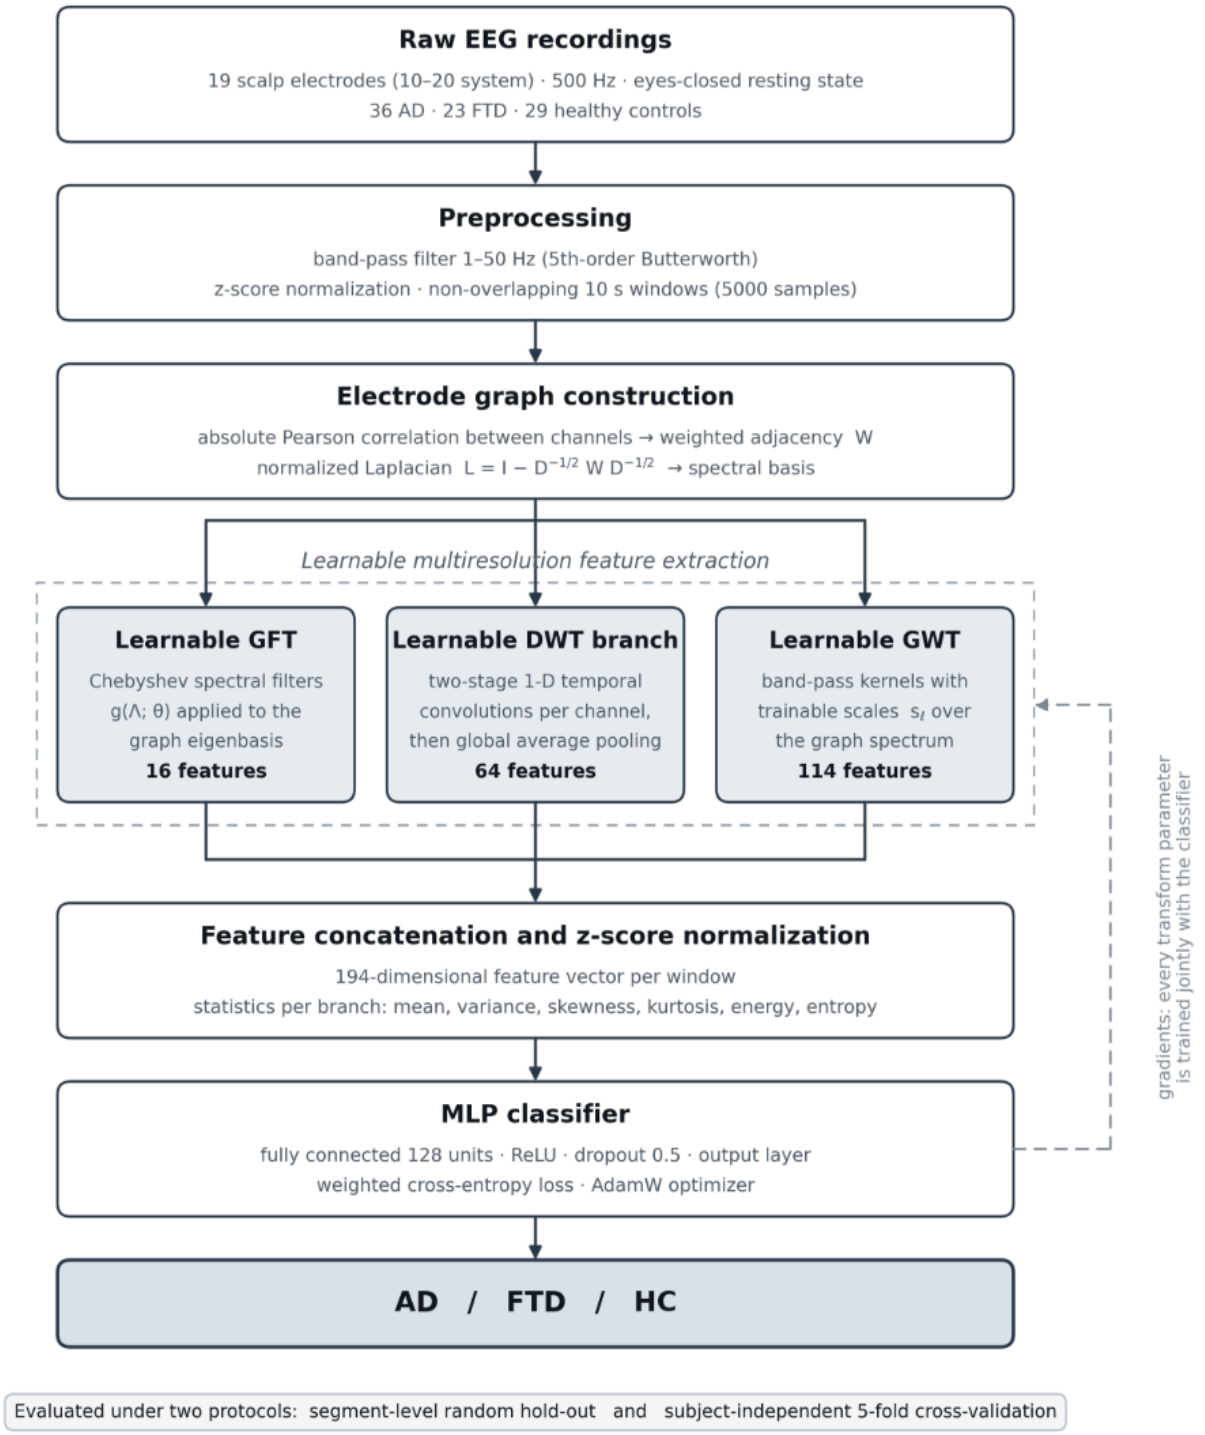


Figure 1

The proposed learnable multiresolution graph signal processing framework. The normalized Laplacian of the electrode graph provides the spectral basis for the learnable GFT and GWT branches, while the learnable DWT branch operates directly on the time domain. The three branches, the concatenation stage and the MLP classifier are optimized jointly

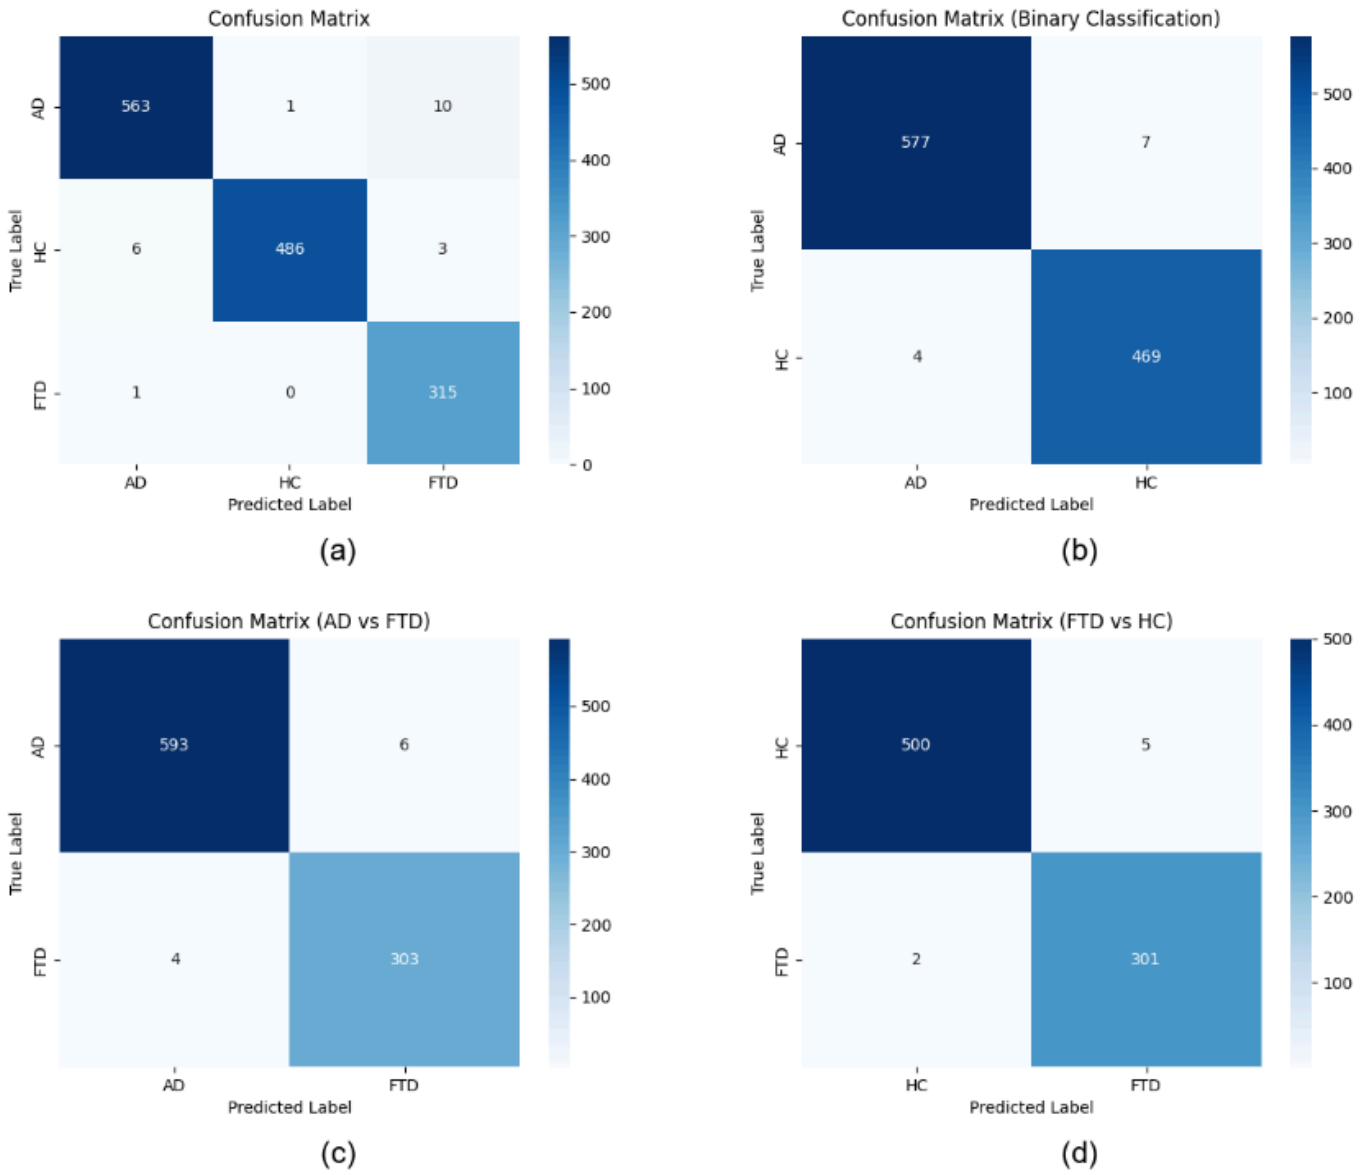


Figure 2

Confusion matrices under the segment-level protocol: (a) AD, FTD and HC, (b) AD vs. HC, (c) AD vs. FTD and (d) FTD vs. HC

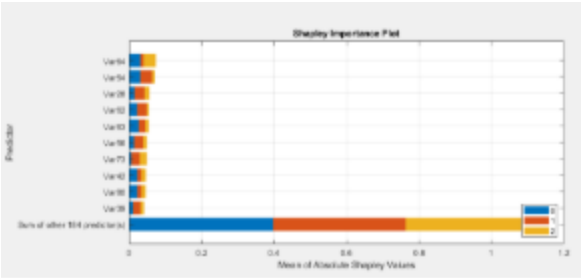


Figure 3

Shapley importance plot: mean absolute Shapley values for each class (0: AD, 1: HC, 2: FTD), aggregated over all predictors