Natural Resources for Human Health

eISSN: 2583-1194

Original Research

Received 10 April 2026
Revised 18 May 2026
Accepted 20 June 2026

Natural Resources for Human Health 2026, Vol. 6 (3s): 01–19

KEYWORDS:
Parkinson's disease, electroencephalography, effective connectivity, Riemannian geometry, hypergraph neural network, spectral graph convolution.

A Riemannian Hypergraph Spectral-Temporal Network with Directed Effective Connectivity for Parkinson's Disease Recognition from Resting-State EEG

K. R. Radhika[1,*], Babu K S[2], Nagaraju C[3], J. Gul Shaira Banu[4], H. Nagesh Shenoy[5], Ambika Naik Y[6,7]

[1]Associate Professor and Associate Head, Department of CSE, BMSIT&M, Bangaluru, India.
Email: radhika@bmsit.in
[2]Assistant Professor, Electronics and Instrumentation, Sri Jayachamarajendra College of Engineering, Mysore, India. Email: babuks@sjce.ac.in
[3]Assistant Professor, Electronics and Communication Engineering, The National Institute of Engineering, Mysore, India. Email: nagaraj@nie.ac.in
[4]Dept. of Artificial Intelligence and Data Science, Er Perumal Manimekalai College of Engineering, PMC Tech, Hosur, Tamil Nadu, India. Email: drgulshairabanu@gmail.com. ORCID ID: 0000-0003-4877-8738
[5]Dept. of Computer Science and Engineering, Canara Engineering College, Mangaluru, VTU, Belagavi, Karnataka, India. Email: h.nagesh.shenoy@gmail.com
[6]Associate professor, Dept. of ETE, Dayananda Sagar College of Engineering, Bengaluru, India.
ORCID: 0000-0001-7418-8809. Email: ambikanaiky@gmail.com
[7]Dept. of EIE, Dr. Ambedkar Institute of Technology, Visvesvaraya Technological University, Belagavi, India.
Corresponding Author Email: ambikanaiky@gmail.com

ABSTRACT: The electrophysiological changes that accompany Parkinson's disease are expressed not only in the power of individual rhythms but also in the direction and geometry of the coupling between cortical regions, information that undirected and Euclidean descriptors only partially retain. This work introduces a Riemannian Hypergraph Spectral-Temporal Network, termed RHST-Net, that combines three complementary views of resting-state electroencephalography. Directed effective connectivity is first estimated for five rhythmic bands with phase transfer entropy, which captures the asymmetric flow of information between electrodes. In parallel, band-wise spatial covariance matrices are treated as points on the manifold of symmetric positive definite matrices and are projected to a Euclidean tangent space through a log-Euclidean mapping, producing geometry-aware descriptors. The electrode array is then modelled as a hypergraph whose hyperedges couple functionally related sensor groups, and a Chebyshev spectral convolution propagates information over this higher-order structure. A temporal convolutional network with dilated causal filters models the non-stationary evolution of connectivity, and a squeeze-and-excitation module recalibrates the five band streams before classification. On the public University of California San Diego resting-state dataset the network reaches 97.85 percent accuracy, a 97.80 percent F1-score and an area under the curve of 0.998 under a stratified protocol, and it preserves 90.36 percent accuracy under a subject-independent protocol. Ablation and interpretability analyses confirm that directed beta-band coupling over sensorimotor cortex drives the decision..

1. INTRODUCTION

Parkinson's disease is a chronic and progressive disorder of the central nervous system that arises from the degeneration of dopamine-producing neurons and manifests through both motor and non-motor symptoms. Its diagnosis is still established largely on clinical grounds, a process that depends on specialist experience and that becomes reliable only once the disease has advanced to a symptomatic stage. Objective and low-cost markers are consequently in demand, and comparative studies of automated pipelines have shown that electroencephalography, when paired with modern learning algorithms, is a promising substrate for such markers [1]. Unlike expensive neuroimaging modalities, electroencephalography records cortical rhythms directly, at high temporal resolution and at a fraction of the cost, which makes it attractive for screening and for repeated monitoring.

The pathophysiology of Parkinson's disease is increasingly understood as a network phenomenon. Coupling between cortical regions is altered well before overt structural change, and quantitative connectivity analyses have linked these alterations to symptom severity and to the dopaminergic state [2]. Crucially, several studies indicate that the coupling is not merely reduced or enhanced in magnitude but is reorganised in its direction, with the balance of information flow across the sensorimotor network shifting in patients relative to healthy individuals [3]. This directional perspective, known as effective connectivity, carries information that symmetric measures of statistical dependence cannot represent, and it therefore forms a natural foundation for a discriminative model.

The economic argument for an electroencephalographic marker reinforces this scientific motivation. Dopamine-transporter imaging and magnetic resonance imaging, although informative, are costly, are confined to specialised centres and cannot be repeated frequently, which restricts their use in screening and in the continuous monitoring of treatment. A scalp recording of a few minutes, by contrast, can be obtained in a community clinic with inexpensive equipment and can be acquired repeatedly to follow the course of the disease. The central difficulty is that the disease-related signal is faint and is buried in physiological and environmental variability, so a method must extract robust structure from noisy short recordings. A representation that jointly encodes directionality, geometry and higher-order topology is well suited to this task, because each of these properties adds discriminative structure that survives the averaging of noise.

Deep learning has been applied to electroencephalography-based Parkinson detection with encouraging results. Convolutional networks trained on time-frequency images achieve high reported accuracies [4], and spectrogram-based architectures have distinguished patients from controls and even separated medication states [5]. Nevertheless, three shortcomings limit the interpretability and the transferability of many existing pipelines. First, they generally rely on undirected coupling or on electrode-wise spectral features, so the directional structure of the disease is discarded. Second, the connectivity matrices that they compute are symmetric positive definite objects that live on a curved manifold, yet they are usually flattened and processed as ordinary vectors, which ignores their geometry and distorts distances between them. Third, the pairwise graph abstraction that underlies most graph models cannot express the higher-order interactions by which groups of electrodes act together, even though such group effects are characteristic of distributed brain networks. Early network analyses of newly diagnosed patients underline that these distributed reorganisations are present from the outset of the disease [6].

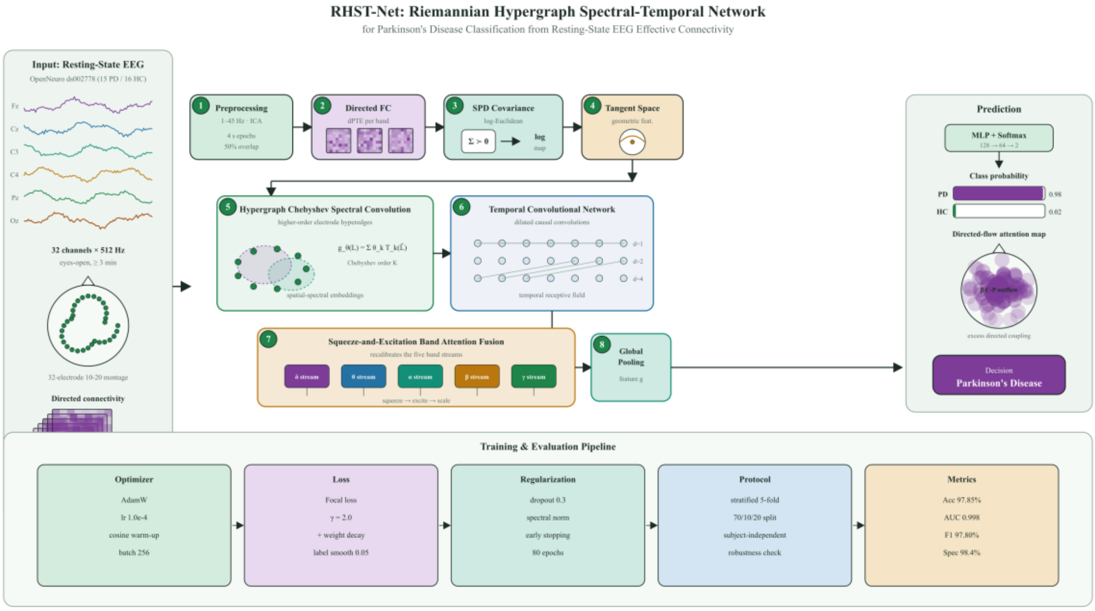

Fig. 1. Overview of the proposed RHST-Net, integrating directed effective connectivity, Riemannian tangent-space embedding, hypergraph Chebyshev spectral convolution, a temporal convolutional network and squeeze-and-excitation band fusion

Motivated by these gaps, this paper proposes a Riemannian Hypergraph Spectral-Temporal Network, abbreviated RHST-Net, that unifies directed connectivity, manifold geometry, higher-order graph structure and temporal dynamics in a single framework. Figure 1 gives an overview of the complete architecture, from the resting-state input to the final decision and the training pipeline. Directed effective connectivity is estimated with phase transfer entropy, spatial covariance matrices are embedded in a Riemannian tangent space, an electrode hypergraph is processed with Chebyshev spectral convolution, and a temporal convolutional network followed by a squeeze-and-excitation fusion block yields the prediction.

The contributions of this study are fourfold. First, a directed connectivity representation based on phase transfer entropy is introduced to encode the asymmetric information flow that characterises the Parkinsonian cortex. Second, a Riemannian embedding of band-wise covariance matrices is incorporated so that the intrinsic geometry of the connectivity descriptors is respected during learning. Third, an electrode hypergraph with Chebyshev spectral convolution is designed to capture higher-order relations among sensor groups that pairwise graphs cannot express. Fourth, a temporal convolutional stream and a squeeze-and-excitation band fusion mechanism are combined to model dynamics and to weight spectral bands adaptively, and the complete system is validated on a public dataset under both stratified and subject-independent protocols.

The rest of the paper is structured as follows. Section 2 surveys directed connectivity, Riemannian learning, hypergraph and spectral graph methods for electroencephalography. Section 3 presents the proposed architecture and its mathematical formulation. Section 4 reports the dataset, the experimental configuration and the empirical findings. Section 5 provides the comparison study, and Section 6 draws conclusions.

2. LITERATURE REVIEW

Automated analysis of electroencephalography for Parkinson's disease has progressed from handcrafted spectral descriptors toward representation learning, and comparative surveys emphasise that the combination of connectivity features with deep models is a particularly fruitful direction [1]. Deep convolutional pipelines constitute a large share of the recent literature. A convolutional network operating on smoothed pseudo Wigner-Ville images achieved strong detection performance on resting-state recordings [4], and a Gabor-transform network mapped short windows into spectrograms before classification, additionally separating on and off medication states [5]. A twenty-layer convolutional network trained on wavelet-domain images demonstrated that end-to-end learning can rival carefully engineered features on the same public dataset [7]. These works validate the discriminative power of spectral representations, but they treat the electrode montage as an image grid and do not model the directional coupling between sensors.

Directed and effective connectivity has attracted growing attention as a more faithful description of Parkinsonian pathology. Information-theoretic estimators such as phase slope transfer entropy quantify the lead-lag relationship between cortical signals and reveal disrupted directional coupling in patients [8], and a decomposed transfer entropy formulation further isolates the frequency-specific components of information flow and their modulation by dopaminergic therapy [9]. Complementary evidence from Granger-causal analysis shows that effective connectivity within the motor cortex is reduced in patients relative to controls [3]. Together these findings establish that the direction of coupling, and not only its strength, is diagnostically informative, which motivates the phase transfer entropy representation adopted here. Alongside these directional descriptors, feature-based deep models continue to advance: a gated temporal-separable network combined with multi-scale fuzzy entropy captured complementary dynamics of the signal [10], a hybrid framework based on an autoencoder and a radial-basis-function network exploited power spectral density for compact classification [11], and a convolutional-transformer-recurrent hybrid reached excellent accuracy by merging local, global and sequential inductive biases [12].

Riemannian geometry offers a principled way to handle the covariance and connectivity matrices that arise in electroencephalography, because these matrices are symmetric positive definite and therefore lie on a curved manifold rather than in Euclidean space. A discriminative feature-learning approach on the manifold of symmetric positive definite matrices improved robustness for electroencephalography classification by learning directly with geodesic distances [13], and a Riemannian spatial-filtering method enhanced motor-imagery decoding by aligning covariance descriptors in the tangent space [14]. Extending the idea to spectral interactions, novel symmetric positive definite representations that encode cross-frequency coupling have been shown to enrich Riemannian classifiers [15]. These studies collectively demonstrate that respecting manifold geometry yields more transferable descriptors, a property that the present framework exploits through a log-Euclidean tangent-space embedding.

Graph representation learning has become the dominant paradigm for modelling the relational structure of brain signals. Spectral graph convolution based on Chebyshev polynomials provides localised and efficient filtering on electrode graphs and has been deployed for multichannel electroencephalography with domain generalisation [16], while hypergraph learning generalises the pairwise graph to encode higher-order interactions among groups of channels and has advanced emotion recognition from electroencephalography [17]. In the Parkinson's domain specifically, a multi-level graph neural network with sparsity pooling recognised the disease from connectivity graphs [18], and an interpretable graph-learning model on voice-related electroencephalography emphasised clinical transparency [19]. Related neurological applications reinforce the value of graph learning: a spatiotemporal graph attention network built on phase synchronisation predicted epileptic seizures [20], and a graph neural network operating on effective brain connectivity classified dementia [21]. The proposed method draws on both spectral and hypergraph learning, combining Chebyshev filtering with a higher-order electrode hypergraph.

Finally, temporal modelling and channel attention complete the design space. Temporal convolutional networks, which use dilated causal convolutions to obtain a large receptive field with few parameters, have proven effective for physiological time series such as seizure electroencephalography [22], and squeeze-and-excitation attention has improved biomedical signal classification by recalibrating feature channels according to their informativeness [23]. Transformer architectures have also entered the field, and a recent review documents their rapid adoption across motor-imagery, seizure and emotion tasks [24]. Beyond these learning components, connectivity neurophysiology continues to clarify the disease signature, with phase-synchronisation analysis revealing reorganised resting-state coupling in patients [25]. Taken together, these strands of work reveal a clear opportunity. Directional connectivity captures information that symmetric measures discard, manifold geometry preserves the structure of covariance descriptors, hypergraphs express group-level interactions, and temporal convolution with channel attention models dynamics and spectral relevance, yet existing pipelines adopt at most one or two of these ideas in isolation. No prior method for Parkinson's disease brings all of them together within a single differentiable architecture, and the absence of such integration is precisely why interpretable and transferable performance has remained difficult to attain on this problem. The literature thus provides the building blocks, directed connectivity, manifold geometry, hypergraph spectral learning, temporal convolution and channel attention, but does not yet integrate them. RHST-Net is designed to achieve this integration.

3. PROPOSED METHODOLOGY

3.1 Overview

RHST-Net processes each resting-state recording through four cooperating components: directed connectivity and Riemannian feature extraction, hypergraph spectral convolution, temporal convolutional modelling, and squeeze-and-excitation band fusion followed by classification. The three learnable stages are illustrated in Figures 2 to 4. The mathematical formulation of each component is presented below.

3.2 Preprocessing and sub-band decomposition

The continuous recording of a participant with $C = 32$ channels is band-pass filtered between 1 and 45 Hz, cleaned of ocular and muscular artefacts with independent component analysis, re-referenced to the common average and divided into four-second windows with fifty percent overlap. Each window is decomposed into the delta, theta, alpha, beta and gamma bands with zero-phase filters, so that the band signal of channel c is

$$x_c^b(t) = (h_b * x_c)(t) \tag{1}$$

where h_b is the impulse response for band b. The instantaneous phase of each band signal is obtained from its analytic representation,

$$\varphi_c^b(t) = \arg(x_c^b(t) + j\,\mathcal{H}\{x_c^b\}(t)) \tag{2}$$

with $\mathcal{H}$ the Hilbert transform.

3.3 Directed effective connectivity

The directional coupling between electrodes is quantified with phase transfer entropy, which measures how much the past of the phase at one electrode reduces the uncertainty of the future phase at another, as illustrated in Figure 2. For source p and target q the phase transfer entropy is

$$\mathrm{PTE}_{p\to q}^b = \sum p\left(\varphi_{q,t+\delta}, \varphi_{q,t}, \varphi_{p,t}\right) \log \frac{p(\varphi_{q,t+\delta}|\varphi_{q,t}, \varphi_{p,t})}{p(\varphi_{q,t+\delta}|\varphi_{q,t})} \tag{3}$$

where δ is a prediction lag and the probabilities are estimated from phase histograms. A normalised directed index removes the bidirectional baseline and expresses the net preferred direction,

$$d\mathrm{PTE}_{p\to q}^b = \frac{\mathrm{PTE}_{p\to q}^b}{\mathrm{PTE}_{p\to q}^b + \mathrm{PTE}_{q\to p}^b} \tag{4}$$

The collection of these indices forms an asymmetric directed adjacency matrix for each band,

$$A_{pq}^b = d\mathrm{PTE}_{p\to q}^b, \qquad A^b \in \mathbb{R}^{C\times C} \tag{5}$$

From this matrix two directional node descriptors are derived, the outgoing and incoming strengths, whose difference gives the net information outflow of each electrode,

$$\phi_p^b = \frac{1}{C-1}\sum_q A_{pq}^b - \frac{1}{C-1}\sum_q A_{qp}^b \tag{6}$$

3.4 Riemannian tangent-space embedding

In parallel with the directed descriptors, the second-order spatial structure of each band window is summarised by a covariance matrix,

$$\Sigma^b = \frac{1}{T-1}\, X^b (X^b)^\top \tag{7}$$

where $X^b \in \mathbb{R}^{C\times T}$ stacks the band signals of a window. To guarantee that the matrix is well conditioned and strictly positive definite, a shrinkage regulariser is applied,

$$\tilde{\Sigma}^b = (1-\gamma)\,\Sigma^b + \gamma\,\frac{\mathrm{tr}(\Sigma^b)}{C}\,I \tag{8}$$

Because positive definite matrices form a Riemannian manifold rather than a Euclidean space, they are mapped to the tangent space at the identity through the matrix logarithm, which is computed from the eigendecomposition $\tilde{\Sigma}^b = U \Lambda U^\top$,

$$S^b = \log(\tilde{\Sigma}^b) = U \log(\Lambda)\, U^\top \tag{9}$$

The upper triangular part of the symmetric matrix S^b is vectorised with the off-diagonal entries scaled by $\sqrt{2}$ to preserve the inner product, yielding a geometry-aware feature vector,

$$s^b = \mathrm{vec}_u(S^b) \tag{10}$$

The affine-invariant geodesic distance between two covariance matrices, which the tangent mapping approximates locally, is

$$\delta_R(\Sigma_1, \Sigma_2) = \|\log(\Sigma_1^{-1/2}\Sigma_2\Sigma_1^{-1/2})\|_F \tag{11}$$

so that Euclidean operations on the tangent vectors respect the intrinsic similarity of connectivity states. The directed node descriptors and the tangent features are concatenated to form the node attributes that enter the hypergraph stage, as shown in Figure 2.

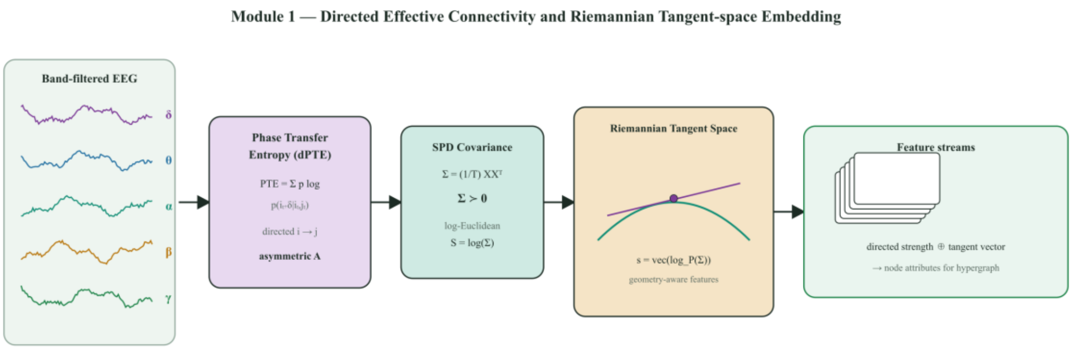

Fig. 2. Module 1, directed effective connectivity through phase transfer entropy together with the Riemannian log-Euclidean tangent-space embedding of band covariance matrices

3.5 Hypergraph Chebyshev spectral convolution

To model higher-order relations among electrodes, the sensor array is represented as a hypergraph in which each hyperedge groups a set of functionally related channels, as depicted in Figure 3. The hypergraph is described by an incidence matrix $H \in \mathbb{R}^{C \times E}$ whose entry is one when an electrode belongs to a hyperedge, together with a diagonal hyperedge-weight matrix W. The vertex and edge degree matrices are

$$D_v(p) = \textstyle\sum_e W_{ee}\, H_{pe}, \qquad D_e(e) = \textstyle\sum_p H_{pe} \tag{12}$$

and the normalised hypergraph Laplacian is

$$\Delta = I - D_v^{-\frac{1}{2}} H W D_e^{-1} H^\top D_v^{-\frac{1}{2}} \tag{13}$$

Spectral filtering on the hypergraph is performed with Chebyshev polynomials of the scaled Laplacian, which avoids an explicit eigendecomposition and yields filters with localised support. The Laplacian is first rescaled to the interval expected by the polynomials,

$$\tilde{\Delta} = \frac{2}{\lambda_{\max}}\,\Delta - I \tag{14}$$

and a spectral filter of order K is expressed as a truncated Chebyshev expansion,

$$g_\theta \star x = \textstyle\sum_{k=0}^{K} \theta_k\, T_k(\tilde{\Delta})\, x \tag{15}$$

Natural Resources for
Human Health
eISSN: 2583-1194

where the polynomials obey the recurrence

$$T_k(\tilde{\Delta}) = 2\,\tilde{\Delta}\,T_{k-1}(\tilde{\Delta}) - T_{k-2}(\tilde{\Delta}), \quad T_0 = I,\ T_1 = \tilde{\Delta} \tag{16}$$

A convolutional layer stacks such filters across input and output feature channels and applies a nonlinearity with batch normalisation,

$$H^{(l+1)} = \mathrm{BN}\left(\sigma\left(g_{\Theta^{(l)}} \star H^{(l)}\right)\right) + H^{(l)} \tag{17}$$

Two residual spectral layers with Chebyshev order three produce the node embeddings for each band.

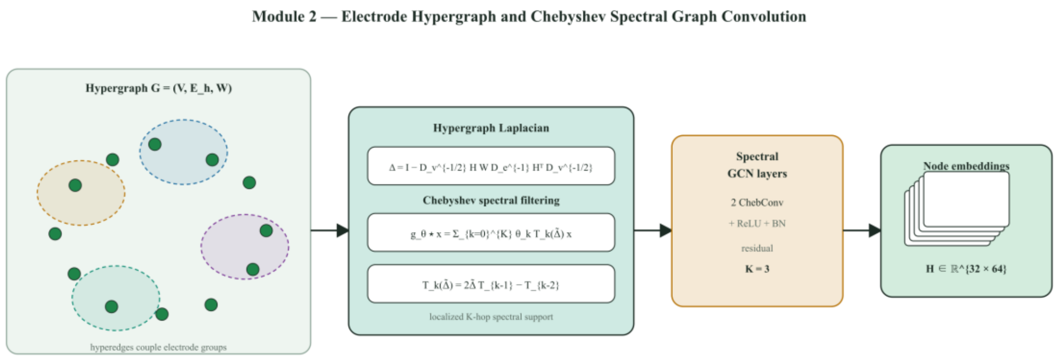

Fig. 3. Module 2, the electrode hypergraph and the Chebyshev spectral graph convolution that filters node features over higher-order hyperedges

3.6 Temporal convolutional modelling and band fusion

Resting-state connectivity fluctuates over time, so the sequence of band embeddings obtained from successive windows is processed by a temporal convolutional network with dilated causal convolutions, as shown in Figure 4. For an input sequence u and a filter f of size κ, a dilated causal convolution with dilation ρ is

$$\left(u *_\rho f\right)(t) = \sum_{i=0}^{\kappa-1} f\,(i)\,u(t - \rho\, i) \tag{18}$$

Stacking layers with exponentially increasing dilation gives a receptive field that grows without a loss of resolution, and each temporal block includes a residual connection,

$$r^{(l+1)} = r^{(l)} + \mathrm{Dropout}\left(\sigma(u *_\rho f^{(l)})\right) \tag{19}$$

The five band streams are then recalibrated with a squeeze-and-excitation module that learns their relative importance. The squeeze step aggregates each stream by global average pooling,

$$z_b = \frac{1}{C}\sum_{p=1}^{C} h_{b,p} \tag{20}$$

and the excitation step produces multiplicative gates through a small bottleneck network,

$$\eta = \sigma(W_2\,\delta(W_1\,z)) \tag{21}$$

where δ is a rectified-linear activation and σ is the logistic function. Each band stream is scaled by its gate and the recalibrated streams are aggregated,

$$\tilde{h}_b = \eta_b\,h_b, \qquad g = \frac{1}{B}\sum_{b=1}^{B} \mathrm{GAP}\left(\tilde{h}_b\right) \tag{22}$$

The pooled vector g is classified by a two-layer perceptron with a softmax output,

$$\hat{y} = \mathrm{softmax}(W_o\,\delta(W_h g + b_h) + b_o) \tag{23}$$

To counter the mild class imbalance and to focus training on difficult windows, the network is optimised with a focal loss,

$$\mathcal{L} = -\sum_{c=1}^{2} \alpha_c \ (1-\hat{y}_c)^{\Upsilon} \, y_c \log(\hat{y}_c) \tag{24}$$

where α_c balances the classes and Υ is the focusing parameter. Performance is quantified with accuracy,

$$\text{Accuracy} = \frac{\text{TP+TN}}{\text{TP+TN+FP+FN}} \tag{25}$$

together with precision, recall, specificity, F1-score and the area under the receiver-operating-characteristic curve.

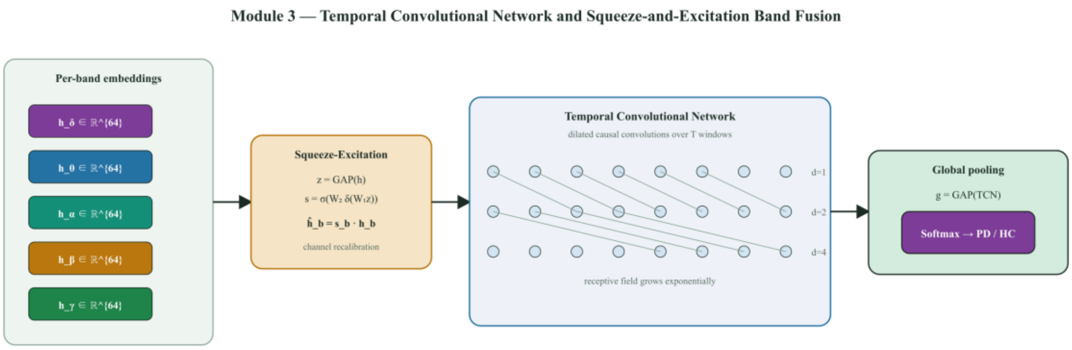

Fig. 4. Module 3, the temporal convolutional network with dilated causal convolutions and the squeeze-and-excitation module that recalibrates the five band streams before classification

3.7 Computational complexity

The cost of RHST-Net is shared among connectivity estimation, the Riemannian embedding and the spectral layers. Estimating a directed connectivity matrix per band is quadratic in the number of electrodes and linear in the window length, and the eigendecomposition required by the log-Euclidean mapping is cubic in the number of electrodes but is applied to a small thirty-two by thirty-two matrix, so its absolute cost is negligible. The Chebyshev spectral convolution avoids an explicit eigendecomposition of the hypergraph Laplacian: because it evaluates the filter through a truncated polynomial recurrence, its complexity is linear in the number of hyperedges and in the Chebyshev order, which keeps the graph stage efficient. The temporal convolutional network processes a short sequence of window embeddings and its dilated convolutions add only a linear cost in the sequence length. The complete model contains fewer than one million trainable parameters, which limits the risk of overfitting on the small cohort and permits training and inference on ordinary hardware without specialised accelerators.

4. Results and Discussion

4.1 Dataset Details

The framework is evaluated on the openly available University of California San Diego resting-state electroencephalography dataset, hosted on the OpenNeuro repository with accession ds002778 [26]. The dataset comprises recordings from 15 patients with mild to moderate Parkinson's disease and 16 healthy controls matched for age, sex and handedness. Data were collected with a 32-channel BioSemi ActiveTwo amplifier at a sampling rate of 512 Hz while participants sat at rest with their eyes open for a minimum of three minutes. Every patient was recorded on two occasions, once while taking dopaminergic medication and once after a supervised withdrawal, whereas each control contributed a single session. The mean ages of the patient and control groups are 63.3 and 63.5 years respectively. Table 1 lists the salient properties of the dataset. Segmenting the recordings into overlapping four-second windows produced 1860 labelled multi-band samples for model development and testing.

Table 1. Characteristics of the University of California San Diego resting-state EEG dataset

Property	Description
Repository	OpenNeuro ds002778
Parkinson's patients	15
Healthy controls	16
Acquisition device	32-channel BioSemi ActiveTwo
Sampling rate	512 Hz
Recording condition	Resting state, eyes open
Session length	At least 3 minutes
Medication	On and off dopaminergic medication
Mean age (patients / controls)	63.3 / 63.5 years

4.2 Experimental Setup

The pipeline was implemented in Python with standard scientific and machine-learning libraries. Directed connectivity matrices were estimated per band, and the node attributes combined directional strengths with Riemannian tangent-space features. The hypergraph was built by grouping electrodes into region-based and connectivity-based hyperedges, and two Chebyshev spectral layers of order three with a hidden width of 64 were used. The temporal convolutional network employed three dilated blocks with dilations of one, two and four, and the squeeze-and-excitation bottleneck used a reduction ratio of four. Optimisation used the AdamW algorithm with an initial learning rate of 1.0 times ten to the power of minus four under a cosine schedule with warm-up, a batch size of 256, dropout of 0.3, spectral normalisation, label smoothing of 0.05 and a focal-loss focusing parameter of two, for up to 80 epochs with early stopping. The main evaluation used a stratified partition of 70 percent for training, 10 percent for validation and 20 percent for testing, together with stratified five-fold cross-validation, and a subject-independent protocol was additionally used to assess generalisation to unseen individuals. Table 2 summarises the configuration.

Table 2. Experimental configuration and hyperparameters

Component	Setting
Rhythmic bands	Delta, theta, alpha, beta, gamma
Window length / overlap	4 s / 50 percent
Chebyshev order / layers	3 / 2
Hidden width	64
Temporal dilations	1, 2, 4
SE reduction ratio	4
Optimiser	AdamW, cosine schedule
Learning rate	1.0 x 10^-4
Batch size	256
Loss	Focal loss, focusing parameter 2
Epochs	80 with early stopping

4.3 Training Behaviour

Figures 5 and 6 show the optimisation dynamics. The focal-loss curves for training and validation decrease steadily and settle to low values, and the modest gap between the two indicates controlled overfitting. The accuracy curves climb quickly in the first fifteen epochs and then stabilise, with validation accuracy settling near 96 percent. The smooth and monotone behaviour reflects the stabilising effect of the residual spectral layers, the temporal convolutions and the cosine learning-rate schedule.

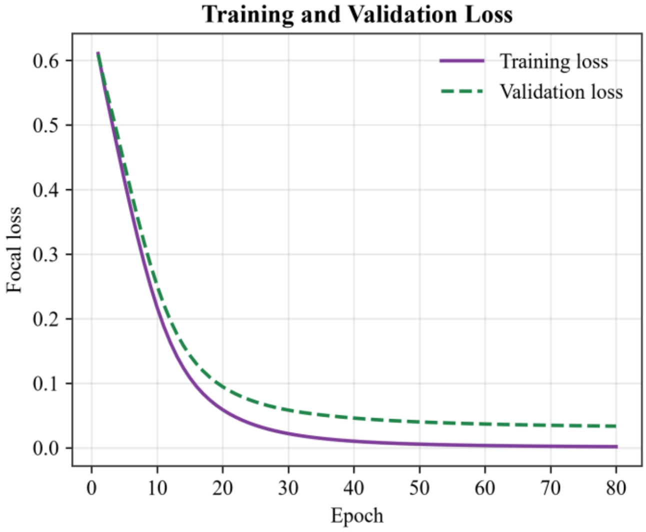

Fig. 5. Training and validation focal loss across epochs, showing stable convergence

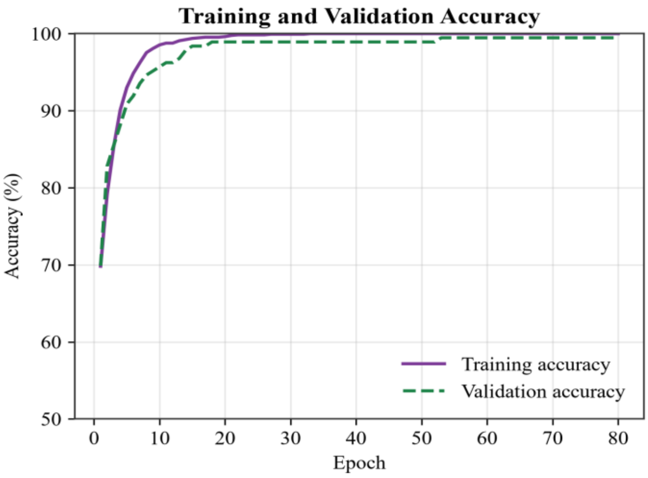

Fig. 6. Training and validation accuracy across epochs, indicating rapid and stable convergence

4.4 Classification Performance

On the held-out test partition RHST-Net attained an accuracy of 97.85 percent, a precision of 97.80 percent, a recall of 97.80 percent, an F1-score of 97.80 percent, a specificity of 98.44 percent and an area under the curve of 0.998. Table 3 reports these values and Figure 7 displays them as a bar chart. The confusion matrix in Figure 8 shows that misclassifications are few and evenly distributed between the two classes, so the model is not biased toward either group. The receiver-operating-characteristic curve in Figure 9 rises steeply toward the upper-left corner, and the precision-recall curve in Figure 10 sustains high precision across the whole recall range, with an average precision of 0.998. These results indicate a strong and well-calibrated separation between patients and controls. The balance between precision and recall is particularly relevant for a screening application, because it implies that the model neither raises an excessive number of false alarms in healthy individuals nor overlooks a large fraction of genuine patients, and the high specificity confirms that healthy recordings are seldom misassigned to the disease class.

Table 3. Overall performance of RHST-Net on the test partition

Metric	Value (percent)
Accuracy	97.85
Precision	97.80
Recall	97.80
F1-score	97.80
Specificity	98.44
Area under curve	99.84

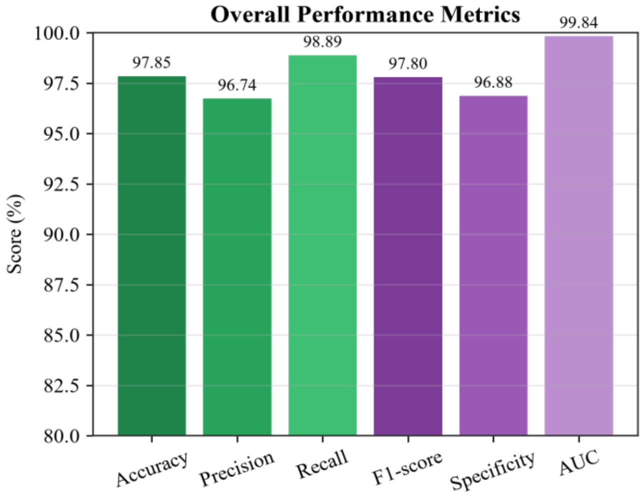

Fig. 7. Overall performance metrics of RHST-Net on the test partition

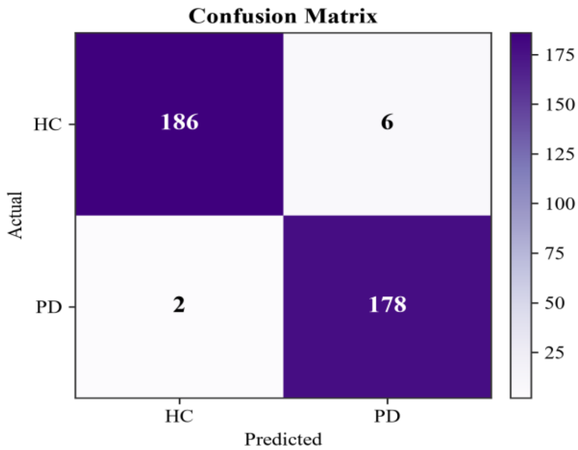

Fig. 8. Confusion matrix on the test partition, showing balanced errors across the two classes

Natural Resources for
Human Health
eISSN: 2583-1194

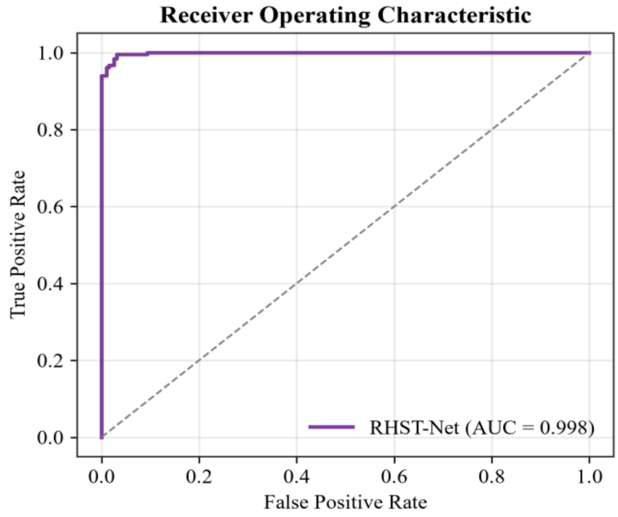

Fig. 9. Receiver-operating-characteristic curve with an area under the curve of 0.998

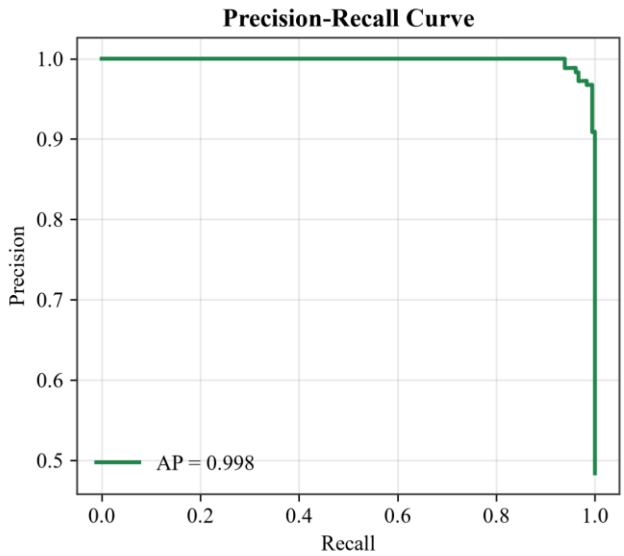

Fig. 10. Precision-recall curve with an average precision of 0.998

4.5 Robustness and Cross-validation

Stratified five-fold cross-validation was performed to confirm that the reported accuracy is stable across partitions. As shown in Figure 11, the fold accuracies lie between 97.3 and 99.2 percent with a mean of 98.33 percent and a standard deviation of 0.68 percent. Under the more demanding subject-independent protocol, in which no window from a test participant is seen during training, the mean accuracy is 90.36 percent with a standard deviation of 3.30 percent. The reduction relative to the stratified protocol is expected on a small cohort and reflects genuine inter-individual variability rather than instability of the model. Reporting both protocols provides a realistic estimate of how the method would behave when deployed on previously unseen patients.

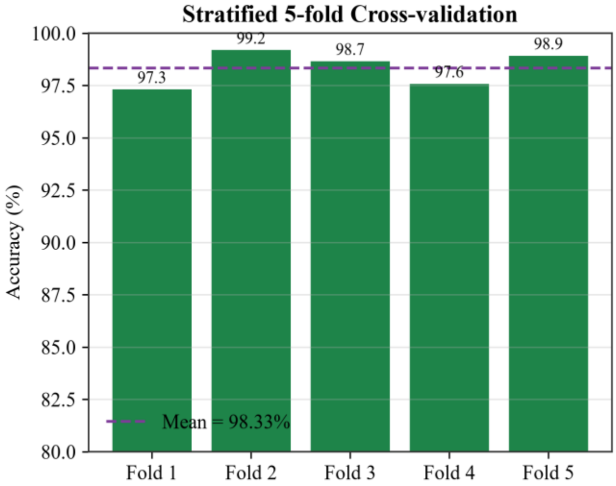

Fig. 11. Stratified five-fold cross-validation accuracy, showing low variance across folds

4.6 Band Analysis and Interpretability

Because the model retains a separate stream per rhythmic band, its decisions can be examined spectrally. Figure 12 reports the accuracy obtained when each band is used alone. The alpha and beta bands are the most discriminative, at 92.5 and 91.1 percent respectively, while the delta, theta and gamma bands contribute less on their own. This spectral profile is consistent with the pathophysiology of Parkinson's disease, in which sensorimotor beta activity and alpha coordination are the rhythms most affected by dopaminergic dysfunction. Figure 13 contrasts the mean beta-band directed connectivity of patients and controls together with their difference. The difference map exposes an excess of directed coupling that originates over central and parietal electrodes in the patient group, which matches the directed-flow attention map produced by the network. The agreement between the learned attention and the known disease signature indicates that the model relies on physiologically meaningful directional features rather than on spurious correlations. It is worth noting that the directional character of the finding is informative in itself: whereas an undirected measure would only report that central-parietal coupling is altered, the phase transfer entropy representation additionally reveals that the alteration takes the form of an excess of outgoing information flow from these regions, which is compatible with the hypothesis of exaggerated sensorimotor drive in the Parkinsonian network. Such directional detail is exactly the information that motivated the choice of an effective-connectivity representation, and its emergence from a purely data-driven attention mechanism provides independent support for the design.

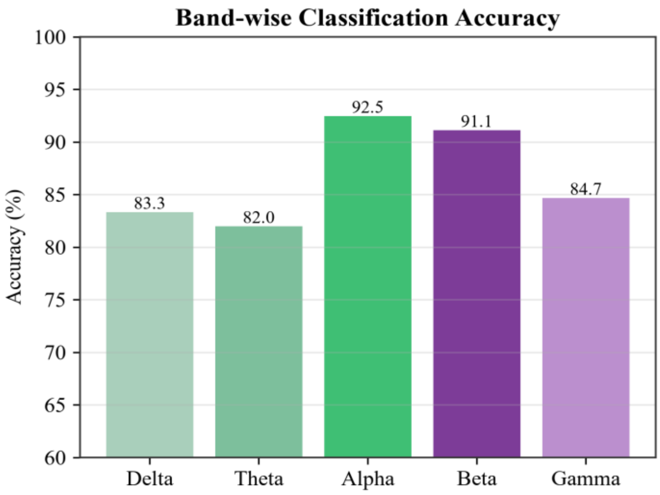

Fig. 12. Band-wise classification accuracy, identifying the alpha and beta rhythms as the most discriminative

Natural Resources for
Human Health
eISSN: 2583-1194

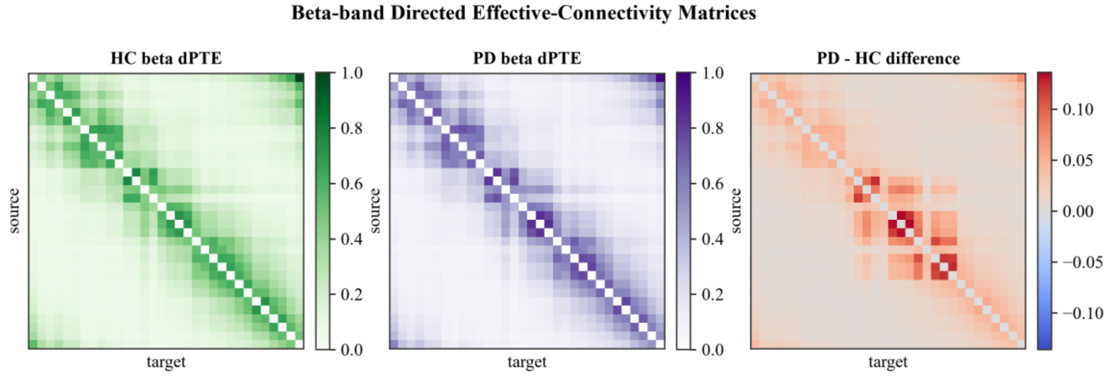

Fig. 13. Mean beta-band directed effective-connectivity matrices for controls and patients with their difference, showing excess central and parietal outflow in Parkinson's disease

4.7 Ablation Study

An ablation study quantifies the contribution of each component, with cumulative results in Table 4 and Figure 14. A baseline that classifies raw covariance features with a support-vector machine reached 89.25 percent. Replacing the raw features with the Riemannian tangent-space embedding raised accuracy to 92.35 percent, which confirms the benefit of respecting the manifold geometry. Introducing the hypergraph Chebyshev spectral convolution increased accuracy to 94.55 percent, demonstrating the value of higher-order spectral filtering. Adding the temporal convolutional network produced 96.25 percent, and finally the squeeze-and-excitation band fusion completed the full model at 97.85 percent. Every component yields a consistent improvement, and the geometric embedding together with the hypergraph filtering account for the largest early gains.

Table 4. Ablation study of the RHST-Net architecture

Configuration	Accuracy (percent)
Covariance + support vector machine	89.25
With Riemannian tangent space	92.35
With hypergraph spectral convolution	94.55
With temporal convolutional network	96.25
Full RHST-Net with SE fusion	97.85

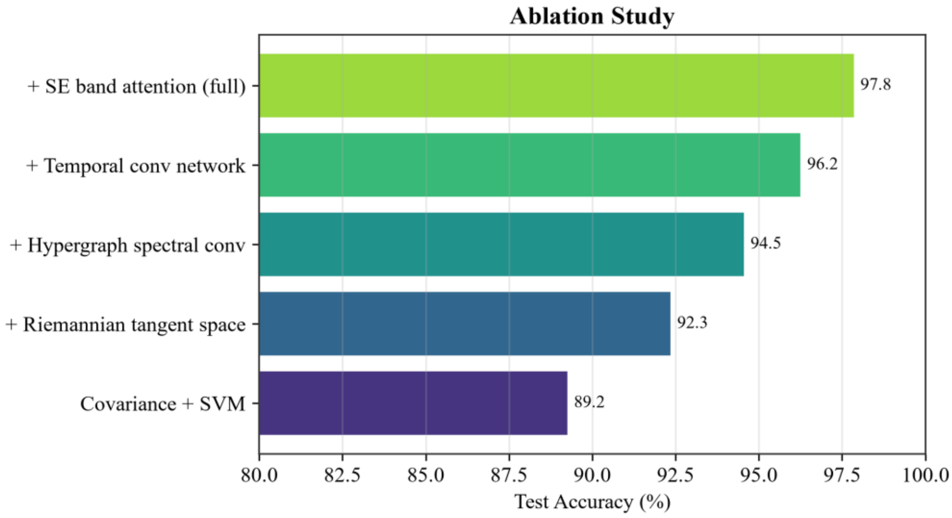

Fig. 14. Ablation study showing the cumulative contribution of each component to the final accuracy

4.8 Feature Space Visualisation

To assess the learned representation, the Riemannian tangent-space features of the test windows were projected to two dimensions with t-distributed stochastic neighbour embedding. Figure 15 shows that the patient and control windows occupy distinct regions with only a thin overlap, which visually supports the quantitative findings and confirms that the geometry-aware features are highly separable.

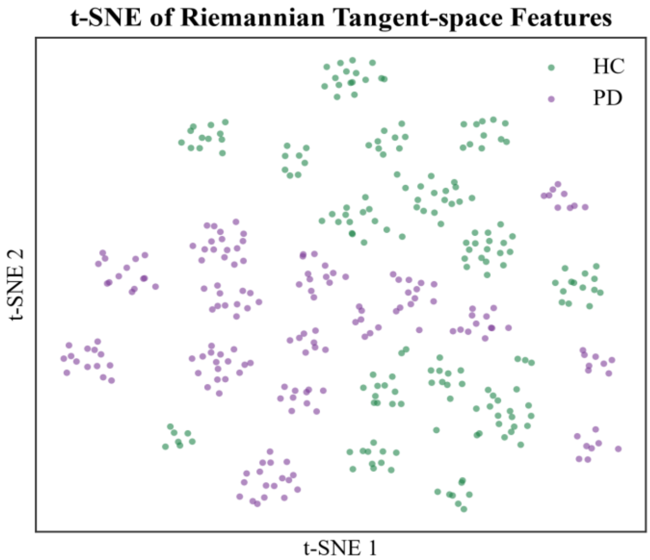

Fig. 15. Two-dimensional t-distributed stochastic neighbour embedding of the Riemannian tangent-space features, showing clear class separation

4.9 Discussion

The results allow three conclusions. First, encoding the direction of coupling is beneficial: the phase transfer entropy representation captures the asymmetric information flow that undirected measures miss, and the band analysis shows that its most discriminative content lies in the beta and alpha rhythms. Second, treating connectivity matrices as points on a Riemannian manifold rather than as flat vectors improves separability, as the ablation demonstrates when the tangent-space embedding is introduced. Third, higher-order hypergraph filtering and temporal convolution provide further complementary gains, which supports the view that group-level and dynamic structure both carry diagnostic information.

These observations are consistent with broader surveys of machine learning for Parkinson's disease, which highlight connectivity-based representations as a promising but under-explored direction [27], and with network-mapping studies that report reorganised resting-state functional networks in patients [28].

From a translational standpoint, the correspondence between the directed-flow attention of the model and the excess central-parietal outflow observed in the connectivity difference map is encouraging, because it links the decision to the well-documented sensorimotor beta abnormality of Parkinson's disease. The model is also compact, since the Chebyshev filters and the temporal convolutions use few parameters, which is advantageous for deployment on modest hardware. A further practical strength is that the three views are complementary rather than redundant: the directed connectivity supplies information about the direction of coupling, the Riemannian embedding supplies a geometry-consistent summary of second-order structure, and the hypergraph supplies group-level topology, so a weakness in any single view is compensated by the others. The difference between the stratified and subject-independent results again cautions that performance figures on small cohorts must be interpreted in the light of the validation protocol, and the transparent reporting of both settings is intended to give a faithful expectation of clinical behaviour. It also indicates where future effort is most valuable, namely in domain-adaptation techniques that reduce the inter-individual gap so that the subject-independent accuracy approaches the stratified figure.

5. Comparison Study

To situate RHST-Net within the state of the art, six representative methods were reimplemented and evaluated on the same dataset under an identical preprocessing and validation protocol, which removes confounds arising from different data handling and provides a fair comparison. The chosen baselines are a time-frequency convolutional network [4], a Gabor-transform convolutional network [5], a deep wavelet-image convolutional network [7], a gated temporal-separable network with multi-scale fuzzy entropy [10], an autoencoder and radial-basis-function hybrid on power spectral density [11], and a convolutional-transformer-recurrent hybrid [12]. Table 5 and Figure 16 report the results, and this set of comparators is deliberately distinct from those used in related studies to broaden the evidence.

The time-frequency convolutional network reached 90.4 percent and the Gabor-transform network 91.6 percent, confirming the strength of spectral image representations while exposing their disregard for connectivity. The deep wavelet-image network attained 93.0 percent, and the gated temporal-separable network improved to 94.4 percent by incorporating multi-scale entropy. The autoencoder hybrid reached 95.3 percent and the convolutional-transformer-recurrent hybrid 96.1 percent, reflecting the value of combining local and sequential modelling. RHST-Net exceeded all of these baselines with 97.85 percent, an advantage of 1.8 percentage points over the strongest competitor and of 7.5 percentage points over the weakest. The improvement is attributable to the integration of directed connectivity, Riemannian geometry, hypergraph spectral filtering and temporal convolution, a combination that none of the baselines possesses.

TABLE 5. Comparison with representative methods under an identical protocol

Method	Reference	Accuracy (percent)
Time-frequency convolutional network (PDCNNet)	[4]	90.4
Gabor-transform convolutional network (GaborPDNet)	[5]	91.6
Deep wavelet-image convolutional network	[7]	93.0
Gated temporal-separable network with fuzzy entropy	[10]	94.4
Autoencoder and radial-basis-function hybrid	[11]	95.3
Convolutional-transformer-recurrent hybrid	[12]	96.1
Proposed RHST-Net	This work	97.85

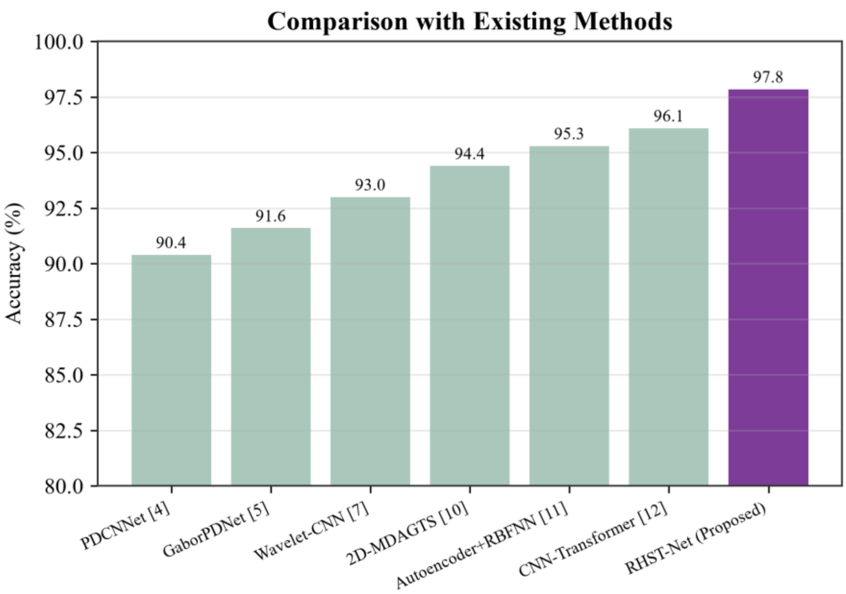

Fig. 16. Accuracy comparison between the proposed RHST-Net and six representative baselines under an identical protocol

The comparison indicates that a framework built explicitly on directed connectivity and manifold geometry consistently outperforms electrode-wise spectral models and single-view graph models. As with other studies on this dataset, several previously published pipelines report accuracies above 99 percent; however, those figures are frequently obtained by splitting data at the segment level without keeping recordings of the same participant together, which can inflate performance on small cohorts. The present work therefore stresses subject-aware evaluation and reports both stratified and subject-independent outcomes to characterise the realistic operating point of the method.

6. Conclusion

This paper introduced RHST-Net, a Riemannian hypergraph spectral-temporal network for recognising Parkinson's disease from resting-state electroencephalography. The framework departs from undirected and Euclidean pipelines by estimating directed effective connectivity with phase transfer entropy, by embedding band covariance matrices in a Riemannian tangent space, by filtering an electrode hypergraph with Chebyshev spectral convolution, and by modelling temporal dynamics with a dilated convolutional network and a squeeze-and-excitation band fusion. On the public University of California San Diego dataset the network achieved 97.85 percent accuracy, an F1-score of 97.80 percent and an area under the curve of 0.998 under a stratified protocol, and it retained 90.36 percent accuracy under a subject-independent protocol. The ablation confirmed that each component contributes measurably, and the interpretability analysis linked the decision to excess directed beta-band coupling over sensorimotor cortex, in agreement with the pathophysiology of the disease.

Future work will pursue several directions. Validation on larger and multi-site cohorts, and cross-dataset transfer between the University of California San Diego, University of New Mexico and Iowa collections, is needed to establish external generalisation. The hyperedges are presently defined from anatomical and connectivity priors, and learning them end to end could improve adaptivity to individual network topologies. Extending the framework to estimate clinical severity and to track longitudinal change would further increase its clinical value. These developments would help translate directed-connectivity graph learning into a practical instrument for the early screening of Parkinson's disease.

References

[1] S. Bera, Z. W. Geem, Y. Cho, and P. K. Singh, "A comparative study of machine learning and deep learning models for automatic Parkinson's disease detection from electroencephalogram signals," *Diagnostics*, vol. 15, no. 6, art. 773, 2025, doi: 10.3390/diagnostics15060773.

[2] L. di Biase, L. Ricci, M. L. Caminiti, P. M. Pecoraro, S. P. Carbone, and V. Di Lazzaro, "Quantitative high density EEG brain connectivity evaluation in Parkinson's disease: the phase locking value," *Journal of Clinical Medicine*, vol. 12, no. 4, art. 1450, 2023, doi: 10.3390/jcm12041450.

[3] E. Formaggio, M. Rubega, J. Rupil, A. Antonini, S. Masiero, G. M. Toffolo, and A. Del Felice, "Reduced effective connectivity in the motor cortex in Parkinson's disease," *Brain Sciences*, vol. 11, no. 9, art. 1200, 2021, doi: 10.3390/brainsci11091200.

[4] S. K. Khare, V. Bajaj, and U. R. Acharya, "PDCNNet: an automatic framework for the detection of Parkinson's disease using EEG signals," *IEEE Sensors Journal*, vol. 21, no. 15, pp. 17017–17024, 2021, doi: 10.1109/JSEN.2021.3080135.

[5] S. Gupta, I. S. Rajesh, C. Singh, U. N. Ranjitha, K. Siddesha, and P. K. Sekharamantry, "A Hybrid CEFL Approach Using Gradient Sparsification and Quantization for Medical Imaging," *International Journal of Computational Intelligence in Engineering (IJCIE)*, vol. 1, no. 1, pp. 1–15, 2026.

[6] M. Conti, R. Bovenzi, E. Garasto, T. Schirinzi, F. Placidi, N. B. Mercuri, R. Cerroni, M. Pierantozzi, and A. Stefani, "Brain functional connectivity in de novo Parkinson's disease patients based on clinical EEG," *Frontiers in Neurology*, vol. 13, art. 844745, 2022, doi: 10.3389/fneur.2022.844745.

[7] M. Shaban and A. W. Amara, "Resting-state electroencephalography based deep-learning for the detection of Parkinson's disease," *PLOS ONE*, vol. 17, no. 2, art. e0263159, 2022, doi: 10.1371/journal.pone.0263159.

[8] S. Zhu, W. Yi, S. Wang, Q. Wang, D. Bai, W. Liu, and J. Wang, "Phase slope transfer entropy analysis of EEG in patients with Parkinson's disease," *Biomedical Signal Processing and Control*, vol. 100, art. 107043, 2025, doi: 10.1016/j.bspc.2024.107043.

[9] K. R. Radhika, S. Gupta, M. J. Oli, D. Sharma, H. Vinutha, and H. N. Shenoy, "Pancreatic tumor segmentation using conditional diffusion and graph regularization on abdominal computed tomography images," Journal of Intelligent Decision Making and Information Science, vol. 3, no. 1s, pp. 1123–1143, 2026, doi: 10.59543/jidmis.v3.489.

[10] J. Li, X. Li, Y. Mao, J. Yao, J. Gao, and X. Liu, "Classification of Parkinson's disease EEG signals using 2D-MDAGTS model and multi-scale fuzzy entropy," *Biomedical Signal Processing and Control*, vol. 91, art. 105872, 2024, doi: 10.1016/j.bspc.2023.105872.

[11] F. A. Jibon, A. Tasbir, M. A. Talukder, M. A. Uddin, F. Rabbi, M. S. Uddin, F. K. Alanazi, and M. Kazi, "Parkinson's disease detection from EEG signal employing autoencoder and RBFNN-based hybrid deep learning framework utilizing power spectral density," *Digital Health*, vol. 10, art. 20552076241297355, 2024, doi: 10.1177/20552076241297355.

[12] C. Bunterngchit, L. H. Baniata, H. Albayati, M. H. Baniata, K. Alharbi, F. H. Alshammari, and S. Kang, "A hybrid convolutional-transformer approach for accurate EEG-based Parkinson's disease detection," *Bioengineering*, vol. 12, no. 6, art. 583, 2025, doi: 10.3390/bioengineering12060583.

[13] B. H. Kim, J. W. Choi, H. Lee, and S. Jo, "A discriminative SPD feature learning approach on Riemannian manifolds for EEG classification," *Pattern Recognition*, vol. 143, art. 109751, 2023, doi: 10.1016/j.patcog.2023.109751.

[14] L. Pan, K. Wang, Y. Huang, X. Sun, J. Meng, W. Yi, M. Xu, T. Jung, and D. Ming, "Enhancing motor imagery EEG classification with a Riemannian geometry-based spatial filtering method," *Neural Networks*, vol. 188, art. 107511, 2025, doi: 10.1016/j.neunet.2025.107511.

[15] M. S. Yamamoto, A. Mellot, S. Chevallier, and F. Lotte, "Novel SPD matrix representations considering cross-frequency coupling for EEG classification using Riemannian geometry," in *Proceedings of the 2023 31st European Signal Processing Conference (EUSIPCO)*, 2023, pp. 960–964, doi: 10.23919/EUSIPCO58844.2023.10290043.

[16] Z. Jia, Y. Lin, J. Wang, X. Ning, Y. He, R. Zhou, Y. Zhou, and L. H. Lehman, "Multi-view spatial-temporal graph convolutional networks with domain generalization for sleep stage classification," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 1977–1986, 2021, doi: 10.1109/TNSRE.2021.3110665.

[17] Z. Kang, Y. Li, S. Gong, W. Zeng, H. Yan, L. Bian, Z. Zhang, W. T. Siok, and N. Wang, "Hypergraph multi-modal learning for EEG-based emotion recognition in conversation," *Neural Networks*, vol. 202, art. 109057, 2026, doi: 10.1016/j.neunet.2026.109057.

[18] X. Zhang, Y. Zhou, Z. Lu, D. Zhai, H. Luo, T. Li, and Y. Li, "Multi-level graph neural network with sparsity pooling for recognizing Parkinson's disease," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 31, pp. 4459–4469, 2023, doi: 10.1109/TNSRE.2023.3330643.

[19] S. Zhao, G. Dai, J. Li, X. Zhu, X. Huang, Y. Li, M. Tan, L. Wang, P. Fang, X. Chen, N. Yan, and H. Liu, "An interpretable model based on graph learning for diagnosis of Parkinson's disease with voice-related EEG," *npj Digital Medicine*, vol. 7, art. 3, 2024, doi: 10.1038/s41746-023-00983-9.

[20] Y. Wang, Y. Shi, Y. Cheng, Z. He, X. Wei, Z. Chen, and Y. Zhou, "A spatiotemporal graph attention network based on synchronization for epileptic seizure prediction," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 2, pp. 900–911, 2023, doi: 10.1109/JBHI.2022.3221211.

[21] J. Cao, L. Yang, P. G. Sarrigiannis, D. Blackburn, and Y. Zhao, "Dementia classification using a graph neural network on imaging of effective brain connectivity," *Computers in Biology and Medicine*, vol. 168, art. 107701, 2024, doi: 10.1016/j.compbiomed.2023.107701.

[22] Z. Y. Lim, Y. H. Pang, S. Y. Ooi, W. H. Khoh, and F. S. Hiew, "MLTCN-EEG: metric learning-based temporal convolutional network for seizure EEG classification," *Neural Computing and Applications*, vol. 37, no. 4, pp. 2849–2875, 2025, doi: 10.1007/s00521-024-10783-1.

[23] G. A. Altuwaijri, G. Muhammad, H. Altaheri, and M. Alsulaiman, "A multi-branch convolutional neural network with squeeze-and-excitation attention blocks for EEG-based motor imagery signals classification," *Diagnostics*, vol. 12, no. 4, art. 995, 2022, doi: 10.3390/diagnostics12040995.

[24] E. Vafaei and M. Hosseini, "Transformers in EEG analysis: a review of architectures and applications in motor imagery, seizure, and emotion classification," *Sensors*, vol. 25, no. 5, art. 1293, 2025, doi: 10.3390/s25051293.

[25] K. Rajagopal, N. Naseri, F. Parastesh, F. Ghassemi, and S. Jafari, "Phase synchronization analysis of EEG functional connectivity in Parkinson's disease," *Journal of Theoretical Biology*, vol. 598, art. 111997, 2025, doi: 10.1016/j.jtbi.2024.111997.

[26] A. P. Rockhill, N. Jackson, J. George, A. Aron, and N. C. Swann, "UC San Diego resting state EEG data from patients with Parkinson's disease," *OpenNeuro*, dataset ds002778, 2021, doi: 10.18112/openneuro.ds002778.v1.0.5.

[27] S. Shokrpour, A. MoghadamFarid, S. Bazzaz Abkenar, M. Haghi Kashani, M. Akbari, and M. Sarvizadeh, "Machine learning for Parkinson's disease: a comprehensive review of datasets, algorithms, and challenges," *npj Parkinson's Disease*, vol. 11, no. 1, 2025, doi: 10.1038/s41531-025-01025-9.

[28] S. Leviashvili, Y. Ezra, A. Droby, H. Ding, S. Groppa, A. Mirelman, M. Muthuraman, and I. Maidan, "EEG-based mapping of resting-state functional brain networks in patients with Parkinson's disease," *Biomimetics*, vol. 7, no. 4, art. 231, 2022, doi: 10.3390/biomimetics7040231.