

Modeling Multivariate Biosignals With Graph Neural Networks and Structured State Space Models

Siyi Tang

Stanford University

SIYITANG@STANFORD.EDU

Jared A. Dunnmon

Stanford University

JDUNNMON@STANFORD.EDU

Liangqiong Qu

University of Hong Kong

LIANGQQU@HKU.HK

Khaled K. Saab

Stanford University

KSAAB@STANFORD.EDU

Tina Baykaner

Stanford University

TINA4@STANFORD.EDU

Christopher Lee-Messer

Stanford University

CLEEMESS@STANFORD.EDU

Daniel L. Rubin

Stanford University

RUBIN@STANFORD.EDU

Abstract

Multivariate biosignals are prevalent in many medical domains, such as electroencephalography, polysomnography, and electrocardiography. Modeling spatiotemporal dependencies in multivariate biosignals is challenging due to (1) long-range temporal dependencies and (2) complex spatial correlations between the electrodes. To address these challenges, we propose representing multivariate biosignals as time-dependent graphs and introduce GRAPHS4MER, a general graph neural network (GNN) architecture that improves performance on biosignal classification tasks by modeling spatiotemporal dependencies in biosignals. Specifically, (1) we leverage the Structured State Space architecture, a state-of-the-art deep sequence model, to capture long-range temporal dependencies in biosignals and (2) we propose a graph structure learning layer in GRAPHS4MER to learn dynamically evolving graph structures in the data. We evaluate our proposed model on three distinct biosignal classification tasks and show that GRAPHS4MER consistently improves over existing models, including (1) seizure detection from electroencephalographic signals, outperforming a previous GNN with self-supervised pre-training by 3.1 points in AUROC; (2) sleep staging from polysomnographic signals, a 4.1

points improvement in macro-F1 score compared to existing sleep staging models; and (3) 12-lead electrocardiogram classification, outperforming previous state-of-the-art models by 2.7 points in macro-F1 score.

Data and Code Availability All the datasets used in our study are publicly available. The Temple University Hospital EEG Seizure Corpus (TUSZ) is publicly available at https://isip.piconepress.com/projects/tuh_eeg/html/downloads.shtml. The DOD-H dataset is publicly available at <https://github.com/Dreem-Organization/dreem-learning-open>. The ICBEB ECG dataset is publicly available at <http://2018.icbeb.org/Challenge.html>. Source code is publicly available at <https://github.com/tsy935/graphs4mer>.

Institutional Review Board (IRB)

Since all the datasets used in this study are publicly available, this study does not require IRB approval.

1. Introduction

Multivariate biosignals are measured by multiple electrodes (sensors) and play critical roles in many medical domains. For example, electroencephalograms (EEGs), which measure brain electrical activity using sensors placed on a person’s scalp, are the most common tests for seizure diagnosis. 12-lead electrocardiograms (ECGs), which assess heart electrical activity using electrodes placed on the surface of a person’s body, are crucial to diagnosing heart rhythm disorders or other disease states of the heart muscle.

Several challenges exist in modeling spatiotemporal dependencies in multivariate biosignals. First, most biosignals are sampled at a high sampling rate, which results in long sequences that can be up to tens of thousands of time steps. Moreover, biosignals often involve long-range temporal correlations (Berthouze et al., 2010; Peng et al., 1995). Prior studies on modeling biosignals often preprocess the raw signals using frequency transformations (Tang et al., 2022b; Asif et al., 2020; Shoeibi et al., 2021; Covert et al., 2019; Guillot et al., 2020; Guillot and Thorey, 2021) or divide the signals into short windows and aggregate model predictions post-hoc (Phan and Mikkelsen, 2022; Pradhan et al., 2022; Strodtzoff et al., 2021). However, such preprocessing steps may discard important information encoded in raw signals, as well as neglect long-range temporal dependencies in the signals. Therefore, a model that is capable of modeling long-range temporal correlations is needed to better capture temporal dependencies in biosignals.

Deep sequence models, including recurrent neural networks (RNNs), convolutional neural networks (CNNs), and Transformers, have specialized variants for handling long sequences (Arjovsky et al., 2016; Erichson et al., 2021; Katharopoulos et al., 2020; Choromanski et al., 2021). However, they struggle to scale to long sequences of tens of thousands of time steps (Tay et al., 2020). Recently, the Structured State Space sequence model (S4) (Gu et al., 2022), a deep sequence model based on the classic state space approach, has outperformed previous state-of-the-art models on challenging long sequence modeling tasks, such as the Long Range Arena benchmark (Tay et al., 2020), raw speech classification (Gu et al., 2022), and audio generation (Goel et al., 2022).

Second, sensors have complex, non-Euclidean spatial correlations (Garcia, 2013; Sazgar and Young, 2019). For example, EEG sensors measure highly correlated yet unique electrical activity from differ-

ent brain regions (Michel and Murray, 2012). Graphs are data structures that can represent complex, non-Euclidean correlations in the data (Chami et al., 2022; Bronstein et al., 2017). Previous works have adopted temporal graph neural networks (GNNs) in modeling multivariate time series, such as EEG-based seizure detection (Covert et al., 2019) and classification (Tang et al., 2022b), traffic forecasting (Li et al., 2018b; Wu et al., 2019; Zheng et al., 2020b; Jiang and Luo, 2022; Tian and Chan, 2021), and pandemic forecasting (Panagopoulos et al., 2021; Kapoor et al., 2020). Nevertheless, most of these studies use sequences up to hundreds of time steps and require a predefined, static graph structure. However, the graph structure of multivariate biosignals may not be easily defined due to unknown sensor locations. For instance, while EEG sensors are typically placed according to the 10-20 standard placement (Jasper, 1958), the exact locations of sensors vary in each individual’s recordings due to the variability in individual head size, clinical need, or patient injury. Moreover, the underlying graph connectivity can evolve over time due to changes in the underlying biology. For example, brain electrical activity is more synchronized across brain areas during the event of a generalized seizure than a normal state (Gloor et al., 1990; Amor et al., 2009), and thus the underlying EEG graph structure can change dynamically when the EEG evolves from the normal state to a seizure state. Similarly, the heart is associated with complex dynamic signals in both normal and diseased states (Qu et al., 2014), and thus the underlying graph structure in ECG can change dynamically over time if a patient develops an abnormal heart rhythm. Hence, the model’s ability to dynamically learn the underlying graph structures is highly desirable.

Graph structure learning (GSL) aims to jointly learn an optimized graph structure and its node and graph representations (Zhu et al., 2021). GSL has been employed for spatiotemporal modeling of traffic flows (Zhang et al., 2020; Tang et al., 2022a; Shang et al., 2021; Wu et al., 2019; Bai et al., 2020), irregularly sampled multivariate time series (Zhang et al., 2022), functional magnetic resonance imaging (fMRI) (El-Gazzar et al., 2021, 2022), and sleep staging (Jia et al., 2020). However, these studies are limited to sequences of less than 1k time steps and do not capture dynamic graphs evolving over time.

In this study, we address the foregoing challenges by (1) leveraging S4 to enable long-range temporal modeling in biosignals, (2) proposing a graph




Figure 1: **Architecture of GraphS4mer.** The model has three main components: (1) stacked S4 layers to learn temporal dependencies in each sensor independently; (2) a graph structure learning (GSL) layer to learn dynamically evolving graph structures; (3) GNN layers to learn spatial dependencies in biosignals based on S4 embeddings and learned graph structures.

structure learning layer to learn dynamically evolving graph structures in multivariate biosignals, and (3) proposing an effective approach to combining S4, GSL, and GNN for modeling multivariate biosignals. Our main contributions are:

- We propose GRAPH S4MER (Figure 1), a general end-to-end GNN for modeling multivariate biosignals. Our modeling contributions are: (1) we leverage S4 to capture *long-range temporal dependencies* in biosignals, (2) our model is able to *dynamically learn the underlying graph structures* in the data without predefined graphs, and (3) our approach is a novel, effective way of combining S4, GSL, and GNN.
- We evaluate GRAPH S4MER on three datasets with distinct data modalities and tasks. Our model consistently outperforms existing methods on (1) seizure detection from EEG signals, outperforming a previous GNN with self-supervised pre-training by 3.1 points in AUROC; (2) sleep staging from polysomnographic signals, outperforming existing sleep staging models by 4.1 points in macro-F1 score; (3) ECG classification, outperforming previous state-of-the-art models by 2.7 points in macro-F1 score.
- Qualitative interpretability analysis suggests that our GSL method learns meaningful graph structures that reflect the underlying seizure classes, sleep stages, and ECG abnormalities.

2. Related Works

Temporal graph neural networks. Temporal GNNs have been widely used for modeling multivariate time series. A recent work by Gao and Ribeiro (2022) shows that existing temporal GNNs

can be grouped into two categories: time-and-graph and time-then-graph. The time-and-graph framework treats a temporal graph as a sequence of graph snapshots, constructs a graph representation for each snapshot, and embeds the temporal evolution using sequence models. For example, GCRN (Seo et al., 2018) uses spectral graph convolution (Defferrard et al., 2016) for node representations and Long Short-Term Memory network (LSTM) (Hochreiter and Schmidhuber, 1997) for temporal dynamics modeling. DCRNN (Li et al., 2018b) integrates diffusion convolution with Gated Recurrent Units (GRU) (Cho et al., 2014) for traffic forecasting. Li et al. (2019) combines relational graph convolution with LSTM for path failure prediction. EvolveGCN (Pareja et al., 2020) treats graph convolutional network (GCN) parameters as recurrent states and uses a recurrent neural network (RNN) to evolve the GCN parameters. In contrast, time-then-graph framework sequentially represents the temporal evolution of node and edge attributes, and uses these temporal representations to build a static graph representation (Xu et al., 2020; Rossi et al., 2020). Gao and Ribeiro (2022) proposes a general framework for time-then-graph, which represents the evolution of node and edge features using two RNNs independently, constructs a static graph from these representations, and encodes them using a GNN. However, most of the existing temporal GNNs are limited to short sequences and adopt RNNs or self-attention for temporal modeling, which may be suboptimal when applied to long sequences. To our knowledge, only one study has leveraged S4 for temporal graphs (El-Gazzar et al., 2022). However, the study is limited to fMRI data and the learned graph is static and shared across all fMRI data.

Graph structure learning. GSL aims to jointly learn a graph structure and its corresponding node (or graph) representations. Briefly, GSL methods fall

into the following three categories (Zhu et al., 2021): (1) metric-based approaches that derive edge weights from node embeddings based on a metric function (Li et al., 2018a; Chen et al., 2020; Zhang and Zitnik, 2020), (2) neural-based approaches that use neural networks to learn the graph structure from node embeddings (Luo et al., 2021; Zheng et al., 2020a; Jia et al., 2020), and (3) direct approaches that treat the adjacency matrix or node embedding dictionary as learnable parameters (Jin et al., 2020; Gao et al., 2020). GSL has been applied to temporal graph data. Zhang et al. (2020) proposes learning graph adjacency matrix from traffic data through learnable parameters. Tang et al. (2022a) uses a multi-layer perceptron and a similarity measure to learn the graph structure for traffic forecasting. Shang et al. (2021) leverages a link predictor (two fully connected layers) to obtain the edge weights. Several works adaptively learn the adjacency matrix from data using learnable node embedding dictionaries (Wu et al., 2019; El-Gazzar et al., 2021, 2022). However, these GSL approaches are limited to short sequences of less than 1k time steps and do not capture dynamically evolving graph structures.

3. Background

3.1. Structured State Space model (S4)

The Structured State Space sequence model (S4) (Gu et al., 2022) is a deep sequence model that is capable of capturing long-range dependencies in long sequences. S4 has outperformed prior state-of-the-art methods on pixel-level 1-D image classification, the challenging Long Range Arena benchmark for long sequence modeling, raw speech classification, and audio generation (Gu et al., 2022; Goel et al., 2022).

S4 is based on state space model (SSM). A recent line of work (Gu et al., 2020, 2021, 2022) has shown that SSMs can be viewed as both CNNs and RNNs. The continuous-time SSM maps an 1-D signal $u(t)$ to a high dimensional latent state $x(t)$ before projecting to an 1-D output signal $y(t)$:

$$x'(t) = \mathbf{A}x(t) + \mathbf{B}u(t), \quad y(t) = \mathbf{C}x(t) + \mathbf{D}u(t) \quad (1)$$

Equation 1 can be used as a black-box sequence representation in a deep learning model, where \mathbf{A} , \mathbf{B} , \mathbf{C} , and \mathbf{D} are parameters learned by gradient descent. In the rest of this section, \mathbf{D} is omitted for exposition because $\mathbf{D}u(t)$ can be viewed as a skip connection and can be easily computed (Gu et al., 2022).

Gu et al. (2021) shows that discrete-time SSM is equivalent to the following convolution:

$$\bar{\mathbf{K}} := (\overline{\mathbf{C}\mathbf{B}}, \overline{\mathbf{C}\mathbf{A}\mathbf{B}}, \dots, \overline{\mathbf{C}\mathbf{A}^{L-1}\mathbf{B}}), \quad y = \bar{\mathbf{K}} * u \quad (2)$$

where $\bar{\mathbf{A}}$, $\bar{\mathbf{B}}$, and $\bar{\mathbf{C}}$ are the discretized matrices of \mathbf{A} , \mathbf{B} , and \mathbf{C} , respectively; $\bar{\mathbf{K}}$ is called the SSM convolution kernel; and L is the sequence length.

S4 (Gu et al., 2022) is a special instantiation of SSMs that parameterizes \mathbf{A} as a diagonal plus low-rank (DPLR) matrix. This parameterization has two major advantages. First, it enables fast computation of the SSM kernel $\bar{\mathbf{K}}$. Second, the parameterization includes the HiPPO matrices (Gu et al., 2020), which are a class of matrices that is capable of capturing long-range dependencies.

However, naively applying S4 for multivariate signals projects the signal channels to the hidden dimension with a linear layer (Gu et al., 2022), which does not take into account the underlying graph structure of multivariate biosignals. This motivates our development of GRAPHs4MER described in Section 4.

3.2. Graph regularization

Graph regularization encourages a learned graph to have several desirable properties, such as smoothness, sparsity, and connectivity (Chen et al., 2020; Kalofolias, 2016; Zhu et al., 2021). Let $\mathbf{X} \in \mathbb{R}^{N \times D}$ be a graph data with N nodes and D features. First, a common assumption in graph signal processing is that features change smoothly between adjacent nodes (Ortega et al., 2018). For instance, in an EEG graph constructed based on the standard EEG placement, neighboring nodes are physically adjacent EEG sensors, and their signals tend to have similar values because these adjacent EEG sensors measure brain electrical activity from adjacent brain areas. Given an undirected graph with adjacency matrix \mathbf{W} , the smoothness of the graph can be measured by the Dirichlet energy (Belkin and Niyogi, 2001):

$$\mathcal{L}_{\text{smooth}}(\mathbf{X}, \mathbf{W}) = \frac{1}{2N^2} \sum_{i,j} \mathbf{W}_{ij} \|\mathbf{x}_{i,:} - \mathbf{x}_{j,:}\|^2 \quad (3)$$

$$= \frac{1}{N^2} \text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}) \quad (4)$$

where $\text{tr}(\cdot)$ denotes the trace of a matrix, $\mathbf{L} = \mathbf{D} - \mathbf{W}$ is the graph Laplacian, \mathbf{D} is the degree matrix of \mathbf{W} . In practice, the normalized graph Laplacian $\hat{\mathbf{L}} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$ is used so that the Dirichlet energy is independent of node degrees. Intuitively, the

eigenvectors of the normalized graph Laplacian are minimizers of the Dirichlet energy (Cai and Wang, 2020), and thus minimizing Equation 4 encourages the learned graph to be smooth.

However, simply minimizing the Dirichlet energy may result in a trivial solution $\mathbf{W} = \mathbf{0}$, i.e., $\mathcal{L}_{\text{smooth}} = 0$ if $\mathbf{W} = \mathbf{0}$. To avoid this trivial solution and encourage sparsity of the graph, additional constraints can be added (Chen et al., 2020; Kalofolias, 2016):

$$\mathcal{L}_{\text{degree}}(\mathbf{W}) = -\frac{1}{N} \mathbf{1}^T \log(\mathbf{W}\mathbf{1}) \quad (5)$$

$$\mathcal{L}_{\text{sparse}}(\mathbf{W}) = \frac{1}{N^2} \|\mathbf{W}\|_F^2 \quad (6)$$

where $\|\cdot\|_F$ is the Frobenius norm of a matrix. Intuitively, $\mathcal{L}_{\text{degree}}$ penalizes disconnected graphs and $\mathcal{L}_{\text{sparse}}$ discourages nodes with high degrees (i.e., encourages the learned graph to be sparse).

4. Methods

Figure 1 illustrates the architecture of GRAPHs4MER. It has three main components: (1) stacked S4 layers, (2) a GSL layer, and (3) GNN layers. Sections 4.1–4.4 describe the details of the problem setup and model architecture.

4.1. Problem setup

To capture non-Euclidean correlations among sensors, we propose a general representation of multivariate biosignals using graphs. Let $\mathbf{X} \in \mathbb{R}^{N \times T \times M}$ be a multivariate biosignal, where N is the number of sensors, T is the sequence length, and M is the input dimension of the signal (typically $M = 1$). We represent the biosignal as a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathbf{W}\}$, where the set of nodes \mathcal{V} corresponds to the sensors (channels/leads), \mathcal{E} is the set of edges, and \mathbf{W} is the adjacency matrix. Here, \mathcal{E} and \mathbf{W} are unknown and will be learned by our model.

While our formulation is general to any type of node or graph classification and regression tasks, we focus on graph classification tasks in this work. Graph classification aims to learn a function which maps an input signal to a prediction, i.e., $f : \mathbf{X} \rightarrow y \in \{1, 2, \dots, C\}$, where C is the number of classes.

4.2. Temporal modeling with S4

As biosignals often involve long-range temporal correlations, we leverage S4 to capture these long-range

temporal dependencies. However, naively applying S4 for multivariate signals projects the N signal channels to the hidden dimension with a linear layer (Gu et al., 2022), which may be suboptimal because it neglects the underlying graph structure of biosignals. Instead, we use stacked S4 layers with residual connection to embed signals in each channel (sensor) independently, resulting in an embedding $\mathbf{H} \in \mathbb{R}^{N \times T \times D}$ (referred to as ‘‘S4 embeddings’’ hereafter) for each input signal \mathbf{X} , where D is the embedding dimension.

4.3. Dynamic graph structure learning

Graph structure learning layer. As motivated in Introduction, the graph structure of multivariate biosignals may not be easily defined due to unknown sensor locations, and the underlying graph connectivity can evolve over time. Hence, we develop a graph structure learning (GSL) layer to learn a *unique* graph structure within a short time interval r , where r is a pre-specified resolution. Instead of learning a unique graph at each time step, we choose to learn a graph over a time interval of length r because (1) aggregating information across a longer time interval can result in less noisy graphs and (2) it is more computationally efficient. For convenience, we let the predefined resolution r be an integer and assume that the sequence length T is divisible by r without overlaps. In the remainder of this paper, we denote $n_d = \frac{T}{r}$ as the number of dynamic graphs.

For graph classification problems, the model’s task is to predict the classes of unseen multivariate biosignals. Therefore, the goal of GSL is to learn a unique graph for each biosignal within r time intervals. Moreover, the GSL layer is expected to be able to generalize to unseen signals. Learning the graph structure is equivalent to learning to compute the edge weight between each pair of nodes. We observe that there is an analogy between self-attention (Vaswani et al., 2017) in natural language processing and GSL—self-attention results in an attention weight between each pair of words in a sentence, and GSL learns the edge weight between each pair of nodes. Given this analogy, we adopt self-attention (Vaswani et al., 2017) to allow each node to attend to all other nodes, and we use the attention weights as the edge weights. The adjacency matrix of the t -th dynamic graph ($t = 1, 2, \dots, n_d$), $\overline{\mathbf{W}}^{(t)} \in \mathbb{R}^{N \times N}$, is

learned as follows:

$$\mathbf{Q} = \mathbf{h}^{(t)}\mathbf{M}^Q, \quad \mathbf{K} = \mathbf{h}^{(t)}\mathbf{M}^K, \quad (7)$$

$$\overline{\mathbf{W}}^{(t)} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{D}}\right), \text{ for } t = 1, 2, \dots, n_d \quad (8)$$

Here, $\mathbf{h}^{(t)} \in \mathbb{R}^{N \times D}$ is mean-pooled S4 embeddings within the t -th time interval of length r ; $\mathbf{M}^Q \in \mathbb{R}^{D \times D}$ and $\mathbf{M}^K \in \mathbb{R}^{D \times D}$ are weights projecting $\mathbf{h}^{(t)}$ to query \mathbf{Q} and key \mathbf{K} , respectively. The above equations can be easily extended to multihead self-attention (Vaswani et al., 2017).

There are two major differences between our GSL layer and prior works (Wu et al., 2019; El-Gazzar et al., 2021, 2022). First, our GSL layer is applied to S4 embeddings rather than raw signals, which leverages the long-range dependencies learned by S4 layers. Second, our GSL layer is able to learn dynamically evolving graph structures over time.

To guide the graph structure learning process, for the t -th dynamic graph, we add a k -nearest neighbor (KNN) graph $\mathbf{W}_{\text{KNN}}^{(t)}$ to the learned adjacency matrix $\overline{\mathbf{W}}^{(t)}$, where each node’s k -nearest neighbors are defined by cosine similarity between their respective values in $\mathbf{h}^{(t)}$ (Chen et al., 2020):

$$\mathbf{W}^{(t)} = \epsilon \mathbf{W}_{\text{KNN}}^{(t)} + (1 - \epsilon) \overline{\mathbf{W}}^{(t)} \quad (9)$$

Here, $\epsilon \in [0, 1)$ is a hyperparameter for the weight of the KNN graph.

To introduce graph sparsity for computational efficiency, we prune the adjacency matrix $\mathbf{W}^{(t)}$ by removing edges whose weights are smaller than a certain threshold (Chen et al., 2020), i.e., $\mathbf{W}_{ij}^{(t)} = 0$ if $\mathbf{W}_{ij}^{(t)} < \kappa$, where κ is a hyperparameter. In this study, we treat the biosignals as undirected, as each pair of sensors is correlated (or anti-correlated) without directionality. To ensure that the learned graphs are undirected, we make the learned adjacency matrix symmetric by taking the average of the edge weights between two nodes.

Graph regularization. To encourage the learned graphs to have desirable graph properties, we include three regularization terms commonly used in the GSL literature (Chen et al., 2020; Kalofolias, 2016; Zhu et al., 2021) (see Section 3.2 and Eqns 4–6). The regularization loss is the weighted sum of the three regularization terms and averaged across all dynamic graphs: $\mathcal{L}_{\text{reg}} = \frac{1}{n_d} \sum_{t=1}^{n_d} \alpha \mathcal{L}_{\text{smooth}}(\mathbf{h}^{(t)}, \mathbf{W}^{(t)}) + \beta \mathcal{L}_{\text{degree}}(\mathbf{W}^{(t)}) + \gamma \mathcal{L}_{\text{sparse}}(\mathbf{W}^{(t)})$, where α , β , and γ are hyperparameters.

4.4. Model architecture

The overall architecture of our model (Figure 1 and Algorithm 1 in Appendix A), GRAPHS4MER, consists of three main components: (1) stacked S4 layers with residual connection to model temporal dependencies in signals within each sensor independently, which maps raw signals $\mathbf{X} \in \mathbb{R}^{N \times T \times M}$ to S4 embedding $\mathbf{H} \in \mathbb{R}^{N \times T \times D}$, (2) a GSL layer to learn dynamically evolving adjacency matrices $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(n_d)}$, and (3) GNN layers to learn spatial dependencies between sensors given the learned graph structures $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(n_d)}$ and node features \mathbf{H} . While our model is general to any kind of GNN layers, we use an expressive GNN architecture—Graph Isomorphism Network (Xu et al., 2019; Hu et al., 2020)—in our experiments.

Finally, a temporal pooling layer and a graph pooling layer are added to aggregate temporal and spatial representations, respectively, which is followed by a fully connected layer to produce a prediction for each multivariate biosignal. The total loss is the sum of the graph regularization loss and the classification loss. Figure 1 shows the overall model architecture, and Algorithm 1 in Appendix A shows the pseudocode for GRAPHS4MER.

5. Experiments

5.1. Experimental setup

In this section, we briefly introduce the datasets and experimental setup in our study. For each experiment, we ran three runs with different random seeds and report mean and standard deviation of the results. See Appendix B–C for details on datasets, baseline models, model training procedures, and hyperparameters.

Seizure detection from EEGs. We first evaluate our model on EEG-based seizure detection. We use the publicly available Temple University Hospital Seizure Detection Corpus (TUSZ) v1.5.2 (Shah et al., 2018), and follow the same experimental setup of seizure detection on 60-s EEG clips as in Tang et al. (2022b). The number of EEG sensors is 19. The sampling rate of each 60-s EEG clip is 200 Hz, resulting in a sequence length of 12k time steps. The task is binary classification of detecting whether or not a 60-s EEG clip contains seizure. The resolution r in GSL layer is chosen as 10-s (i.e., 2,000 time steps), which is inspired by how trained EEG readers analyze EEGs. Following prior studies, we use AUROC,

F1-score, area under precision-recall curve (AUPRC), sensitivity, and specificity as the evaluation metrics.

Sleep staging from polysomnographic signals. Next, we evaluate our model on sleep staging from polysomnographic (PSG) signals. PSG is used in the diagnosis of sleep disorders such as obstructive sleep apnea. We use the publicly available Dreem Open Dataset-Healthy (DOD-H) (Guillot et al., 2020). The number of PSG sensors is 16. The sampling rate is 250 Hz, and thus each 30-s signal has 7,500 time steps. The task is to classify each 30-s PSG signal as one of the five sleep stages: wake, rapid eye movement (REM), non-REM sleep stages N1, N2, and N3. The resolution r is 10-s (i.e., 2.5k time steps). Similar to a prior study Guillot and Thorey (2021), we use macro-F1 score and Cohen’s Kappa as the evaluation metrics.

ECG classification. Lastly, we evaluate our model on a 12-lead ECG classification task. We use the publicly available ICBEb ECG dataset (Liu et al., 2018), and follow the same data split as described in Strodtzoff et al. (2021). The number of ECG channels is 12, with a sampling rate of 100 Hz. The ECG lengths vary from 6-s to 60-s, resulting in variable sequence lengths between 600 and 6k time steps. To allow training in mini-batches across variable sequence lengths, we pad short sequences with 0s and mask out the padded values so that the padded values are not seen by the models. There are 9 classes in the dataset (see Table 8 in Appendix B), including one normal class and 8 abnormal classes indicating diseases. Each ECG is labeled with at least one of the 9 classes; therefore, the task is a multilabel classification task. The resolution r in the GSL layer is set as the actual sequence length of each ECG to allow training in batches. Following prior studies (Liu et al., 2018; Strodtzoff et al., 2021), we use macro-F1, macro-F2, macro-G2, and macro-AUROC as the evaluation metrics. See Appendix D for detailed definitions of these metrics.

5.2. Experimental results

Seizure detection from EEG signals. Table 1 shows our model performance on EEG-based seizure detection and its comparison to existing models. GRAPHs4MER outperforms the previous state-of-the-art, Dist-DCRNN with pre-training, by 3.1 points in AUROC. Notably, Dist-DCRNN was pre-trained using a self-supervised task (Tang et al., 2022b),

whereas GRAPHs4MER was trained from scratch without the need of pre-training.

Sleep staging from PSG signals. Table 2 compares our model performance to existing sleep staging models on DOD-H dataset. GRAPHs4MER outperforms RobustSleepNet, a specialized sleep staging model, by 4.1 points in macro-F1. Note that the baselines preprocess the PSG signals using short-time Fourier transform, whereas our model directly takes in raw PSG signals without the need of preprocessing.

ECG classification. Table 3 compares GRAPHs4MER to existing ECG classification models on ICBEb dataset. GRAPHs4MER provides 2.7 points improvement in macro-F1 score compared to XResNet1D (He et al., 2019), a prior state-of-the-art CNN specialized in time series modeling. Note that all the baselines are trained on 2.5-s ECG windows (i.e., 250 time steps) and the window-wise predictions are aggregated post-hoc to obtain the predictions for each ECG, whereas our model directly takes the entire ECG signal (ranging from 600 to 6k time steps) without the need of sliding windows, which allows modeling long-range temporal dependencies in the entire ECG. Table 9 in Appendix E shows model performance in individual ECG classes. GRAPHs4MER outperforms the baselines on 6 out of 9 classes.

Ablations. We perform ablation studies to investigate the importance of (1) graph-based representation, where we remove the GSL and GNN layers in our model; (2) S4 encoder, where we replace the S4 layers with GRUs (Cho et al., 2014); (3) long-range temporal modeling for GSL and graph representation learning, where we apply GSL and GNN layers to raw signals, followed by S4 layers; and (4) GSL, where we remove the GSL layer and use predefined graphs. For seizure detection on TUSZ dataset, we use the predefined distance-based EEG graph as in Tang et al. (2022b). DOD-H and ICBEb datasets have no predefined graph available, and thus we use a KNN graph obtained based on cosine similarity between raw signals in a pair of nodes (setting K the same as GRAPHs4MER). Table 4 shows the ablation results. There are several important observations. First, GRAPHs4MER consistently outperforms S4, suggesting the effectiveness of representing multivariate biosignals as graphs. Second, when S4 layers are replaced with GRUs, the model performance drops by a large margin. This indicates the effectiveness of S4 in modeling long-range temporal dependencies. Third, applying GSL and GNN layers before S4 layers (“GSL-GNN-S4”) results in lower perfor-

Table 1: **Results of EEG-based seizure detection on TUSZ dataset.** Baseline model results are cited from Tang et al. (2022b). Best and second best results are **bolded** and underlined, respectively.

Model	AUROC	F1-Score	AUPRC	Sensitivity	Specificity
LSTM (Hochreiter and Schmidhuber, 1997)	0.715 \pm 0.016	0.365 \pm 0.009	0.287 \pm 0.026	0.463 \pm 0.060	0.814 \pm 0.053
Dense-CNN (Saab et al., 2020)	0.796 \pm 0.014	0.404 \pm 0.022	0.399 \pm 0.017	0.451 \pm 0.134	0.869 \pm 0.071
CNN-LSTM (Ahmedt-Aristizabal et al., 2020)	0.682 \pm 0.003	0.330 \pm 0.016	0.276 \pm 0.009	0.363 \pm 0.044	0.857 \pm 0.023
Dist-DCRNN w/o Pre-training (Tang et al., 2022b)	0.793 \pm 0.022	0.341 \pm 0.170	0.418 \pm 0.046	0.326 \pm 0.183	<u>0.932 \pm 0.058</u>
Corr-DCRNN w/o Pre-training (Tang et al., 2022b)	0.804 \pm 0.015	0.448 \pm 0.029	0.440 \pm 0.021	0.457 \pm 0.058	0.900 \pm 0.028
Dist-DCRNN w/ Pre-training (Tang et al., 2022b)	<u>0.875 \pm 0.016</u>	<u>0.571 \pm 0.029</u>	<u>0.593 \pm 0.031</u>	<u>0.570 \pm 0.047</u>	0.927 \pm 0.012
Corr-DCRNN w/ Pre-training (Tang et al., 2022b)	0.850 \pm 0.014	0.514 \pm 0.028	0.539 \pm 0.024	0.502 \pm 0.047	0.923 \pm 0.008
GRAPHS4MER (ours)	0.906 \pm 0.012	0.680 \pm 0.012	0.723 \pm 0.023	0.718 \pm 0.041	0.933 \pm 0.009

 Table 2: **Results of sleep staging on DOD-H dataset.** Best and second best results are **bolded** and underlined, respectively.

Model	Macro-F1	Kappa
LSTM (Hochreiter and Schmidhuber, 1997)	0.609 \pm 0.034	0.539 \pm 0.046
SimpleSleepNet (Guillot et al., 2020)	0.720 \pm 0.001	0.703 \pm 0.013
RobustSleepNet (Guillot and Thorey, 2021)	<u>0.777 \pm 0.007</u>	<u>0.758 \pm 0.008</u>
DeepSleepNet (Supratak et al., 2017)	0.716 \pm 0.025	0.711 \pm 0.032
GRAPHS4MER (ours)	0.818 \pm 0.008	0.802 \pm 0.014

 Table 3: **Results of ECG classification on ICBEB dataset.** Best and second best results are **bolded** and underlined, respectively.

Model	Macro F1-Score	Macro F2-Score	Macro G2-Score	Macro AUROC
InceptionTime (Ismail Fawaz et al., 2020)	0.778 \pm 0.020	0.792 \pm 0.023	0.576 \pm 0.027	0.964 \pm 0.008
XResNet1D (He et al., 2019)	<u>0.782 \pm 0.016</u>	<u>0.803 \pm 0.016</u>	<u>0.587 \pm 0.019</u>	<u>0.971 \pm 0.001</u>
ResNet1D (Wang et al., 2017)	0.772 \pm 0.015	0.788 \pm 0.008	0.570 \pm 0.014	0.963 \pm 0.004
FCN (Wang et al., 2017)	0.736 \pm 0.015	0.764 \pm 0.014	0.538 \pm 0.014	0.950 \pm 0.003
LSTM (Hochreiter and Schmidhuber, 1997)	0.747 \pm 0.018	0.768 \pm 0.018	0.542 \pm 0.018	0.942 \pm 0.004
Bidir-LSTM (Hochreiter and Schmidhuber, 1997)	0.748 \pm 0.009	0.769 \pm 0.004	0.548 \pm 0.005	0.945 \pm 0.002
WaveletNN (Strodthoff et al., 2021)	0.621 \pm 0.013	0.643 \pm 0.019	0.396 \pm 0.014	0.911 \pm 0.002
GRAPHS4MER (ours)	0.809 \pm 0.004	0.804 \pm 0.004	0.609 \pm 0.005	0.977 \pm 0.001

mance than GRAPHS4MER on all datasets, suggesting the effectiveness of capturing long-range temporal dependencies for GSL and graph representation learning. Finally, GRAPHS4MER with GSL outperforms GRAPHS4MER without GSL (with predefined graph) for sleep staging on the DOD-H dataset, and performs marginally better than GRAPHS4MER without GSL for seizure detection on the TUSZ dataset and ECG classification on the ICBEB dataset. Nevertheless, a major advantage of our GSL method is that it does not require prior knowledge about the graph structure, which would be particularly useful when sensor locations are not available. Table 10 in Appendix F shows additional evaluation metrics for the ablation experiments.

Effect of temporal resolution. We examine the effect of temporal resolution r in Figure 3 in Appendix G for seizure detection and sleep staging tasks. For ECG classification, the temporal resolution r is set as the actual sequence lengths due to variable lengths of the ECG signals, and thus is excluded for this analysis. We observe that a smaller value of r tends to result in higher performance than a larger value of r , suggesting the effectiveness of learning dynamically varying graphs.

Effect of KNN graph weight and graph regularization. In Figures 4–5 in Appendix H, we examine the effect of KNN graph weight (i.e., ϵ in Eqn 9) and graph regularization on model performance.

Table 4: **Ablation results.** For S4, we remove the GSL and GNN layers. For GRAPHs4MER w/o S4, we replace S4 with GRUs. For GSL-GNN-S4, we apply GSL and GNN layers to raw signals, followed by S4 layers. For GRAPHs4MER w/o GSL, we remove the GSL layer and use predefined graphs. P-values are computed using bootstrapping with 5,000 replicates with replacement on the run with median performance. Best and second best results are **bolded** and underlined, respectively.

Model	TUSZ		DOD-H		ICBEB	
	AUROC	p-value	Macro-F1	p-value	Macro-F1	p-value
S4	0.824 ± 0.011	0.0002	0.778 ± 0.009	0.0004	0.781 ± 0.003	0.0132
GRAPHs4MER w/o S4	0.705 ± 0.095	0.0002	0.634 ± 0.061	0.0002	0.197 ± 0.005	0.0002
GSL-GNN-S4	0.882 ± 0.014	0.0004	<u>0.797 ± 0.011</u>	0.0164	0.772 ± 0.012	0.0060
GRAPHs4MER w/o GSL	<u>0.899 ± 0.010</u>	0.7864	0.765 ± 0.016	0.0002	<u>0.797 ± 0.012</u>	0.1268
GRAPHs4MER (ours)	0.906 ± 0.012	ref	0.818 ± 0.008	ref	0.809 ± 0.004	ref

Larger KNN graph weights and graph regularization result in marginally better performance.

Computational costs of GSL. In Appendix I and Table 11, we examine the computational costs of GRAPHs4MER across different values of the time interval r in the GSL layer (see Section 4.3). As the time interval r decreases (i.e., number of dynamic graphs increases), the number of model parameters remains the same, whereas the number of multiply-accumulate (MACs) operations increases. This is expected given that the same GSL layer is applied multiple times to construct dynamic graphs.

Interpretation of graphs. To investigate if the learned graphs are meaningful, we visualize mean adjacency matrices for EEG, PSG, and ECG signals in the correctly predicted test samples in Figure 2, grouped by seizure classes, sleep stages, and ECG classes, respectively. To quantify the differences between two mean adjacency matrices, we report the mean and standard deviation of the absolute values of the differences in Figure 2 (referred to as δ). These adjacency matrices were reviewed and interpreted by clinical experts in terms of whether the differences in adjacency matrices between diseased (or sleep) and normal (or wake) classes reflect characteristics of the diseased (or sleep) states, and several clinically-meaningful patterns were found.

For EEG, the magnitude of differences between generalized seizure and non-seizure adjacency matrices (Figure 2v; $\delta = 0.013 \pm 0.020$) is larger than the magnitude of differences between focal seizure and non-seizure adjacency matrices (Figure 2iv; $\delta = 0.008 \pm 0.010$). This suggests that the abnormalities in generalized seizures are more synchronized among the channels, resulting in high correlations. This is consistent with the literature that generalized

seizures are characterized with abnormally synchronized brain activity (Gloor et al., 1990; Amor et al., 2009).

For PSG, we observe that the magnitude of differences between N3 and wake (Figure 2x; $\delta = 0.044 \pm 0.056$) is larger than the magnitude of differences between N1 and wake (Figure 2ix; $\delta = 0.028 \pm 0.031$). This pattern is expected given that N1 is the earliest sleep stage, whereas N3 is the deep sleep stage and is associated with slow brain waves that are not present in other sleep stages (Iber et al., 2007). Mean adjacency matrices for sleep stages N2 and REM are shown in Figure 6 in Appendix J.

For ECG, we find that the magnitude of differences between adjacency matrices of first-degree atrioventricular block (I-AVB) and normal ECG is small (Figure 2xiv; $\delta = 0.014 \pm 0.019$), which is expected as I-AVB abnormality involves only small changes in the ECG (a subtle increase in the PR interval of ECG, not involving any morphologic changes in the p-waves or the QRS complexes). In contrast, the magnitude of differences between adjacency matrices of left bundle branch block (LBBB) and normal ECG is larger, particularly in ECG leads V1-V6 (Figure 2xv; $\delta = 0.033 \pm 0.046$). This finding is clinically meaningful as ECG signals with LBBB demonstrate a pronounced abnormality in the QRS complexes of the ECG, especially in leads V1-V6. In Appendix K, we show example normal ECG, an ECG with I-AVB, and an ECG with LBBB in Figure 7, as well as adjacency matrices for nine ECG classes in Figure 8.

6. Discussion and Limitations

In this study, we propose a novel GSL method to learn dynamically evolving graph structures in multi-

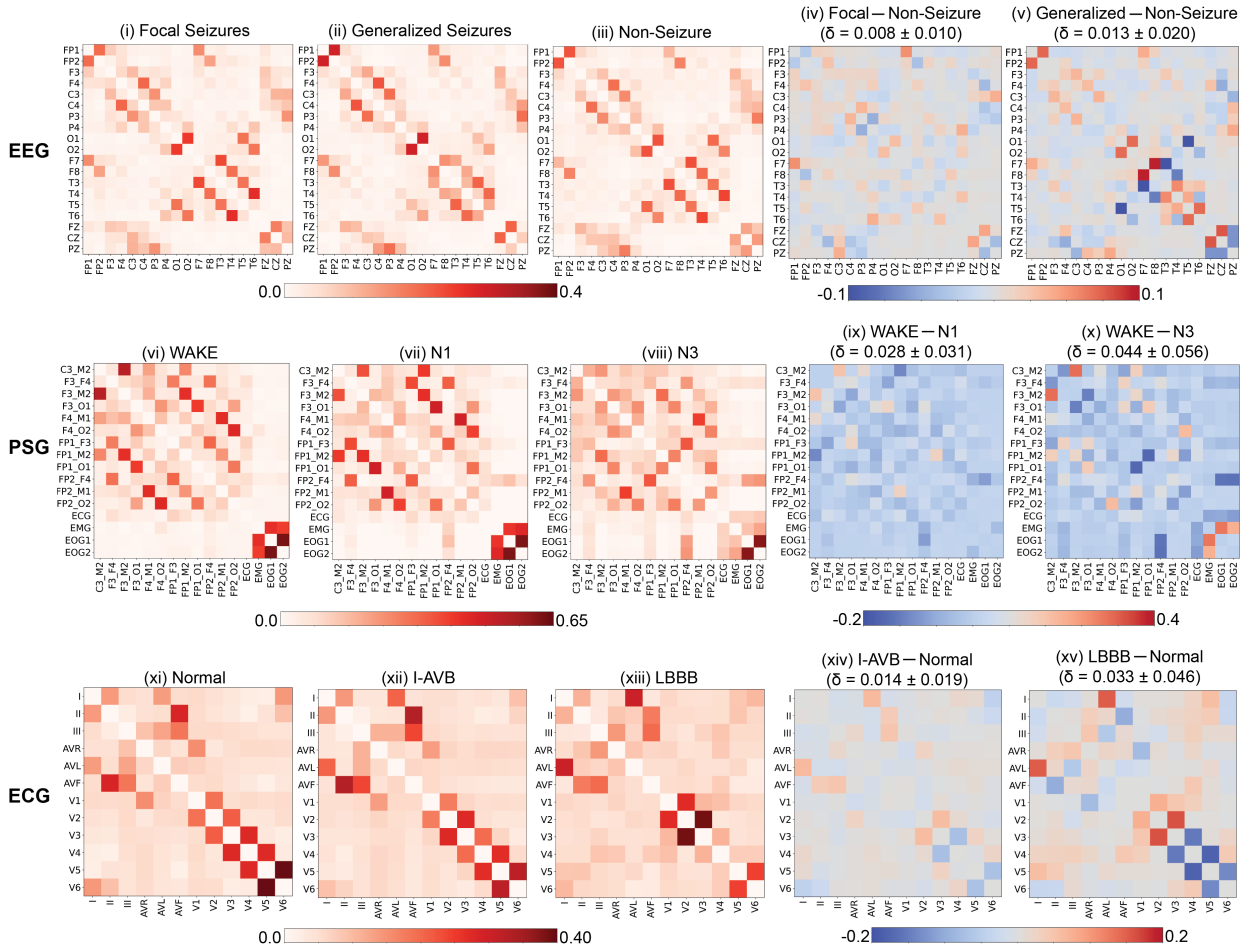


Figure 2: **Top:** Mean adjacency matrix for EEG in correctly predicted test samples for (i) focal seizures, (ii) generalized seizures, and (iii) non-seizure. (iv) Difference between focal seizure and non-seizure adjacency matrices. (v) Difference between generalized seizure and non-seizure adjacency matrices. **Middle:** Mean adjacency matrix for PSG in correctly predicted test samples for sleep stage (vi) WAKE, (vii) N1, and (viii) N3. (ix) Difference between WAKE and N1 adjacency matrices. (x) Difference between WAKE and N3 adjacency matrices. **Bottom:** Mean adjacency matrix for ECG in correctly predicted test samples for (xi) normal ECG, (xii) first-degree atrioventricular block (I-AVB), and (xiii) left bundle branch block (LBBB). (xiv) Difference between I-AVB and normal adjacency matrices. (xv) Difference between LBBB and normal adjacency matrices. δ : mean and std of absolute values of differences between two mean adjacency matrices.

variate biosignals, and present GRAPHs4MER, a general GNN integrating S4 and GSL for modeling multivariate biosignals. Our method sets new state-of-the-art performance on seizure detection, sleep staging, and ECG classification, and learns meaningful graph structures that reflect seizure classes, sleep stages, and ECG abnormalities.

In general, GRAPHs4MER is best suited for biosignals that (1) are measured by multiple sensors placed in or on the surface of a person’s body and (2) exhibit long-range temporal dependencies. The three data types—EEG, PSG, and ECG—satisfy these properties (Berthouze et al., 2010; Peng et al., 1995; Garcia, 2013; Sazgar and Young, 2019), and GRAPHs4MER

outperforms prior state-of-the-art methods on the related classification tasks.

Ablation experiments (Table 4) suggest the importance of each component of GRAPHs4MER. First, GRAPHs4MER consistently outperforms S4, indicating the effectiveness of representing biosignals as graphs. Second, replacing S4 with GRU decreases the model performance by a large margin, which suggests the importance of S4 in modeling long-range temporal dependencies. Notably, the three tasks in our study cover a wide range of sequence lengths ranging from 600 to 12k time steps. Third, applying GSL and GNN layers prior to S4 layers underperforms GRAPHs4MER, which indicates that long-range temporal dependencies learned by S4 are crucial for learning graph structures and graph representations. Lastly, GRAPHs4MER outperforms GRAPHs4MER without GSL (with predefined graphs) on sleep staging, which suggests that GSL could be particularly useful for applications where graph structures cannot be easily defined.

There exists several limitations in this study. First, the GSL layer does not take into account domain knowledge. For example, GSL does not consider the standard EEG placement, and it only performs marginally better than the predefined EEG graph (Table 4). Future work can leverage domain knowledge, such as standard sensor placement, to further improve the model performance. Second, self-attention requires time and memory that grows quadratically with the number of nodes, and thus could be suboptimal when the number of sensors is large. Future work can leverage more efficient attention mechanisms (Child et al., 2019; Kitaev et al., 2020; Dao et al., 2022) to improve the computational efficiency of GSL. Third, due to variable ECG lengths in the ICBEb dataset, we set the temporal resolution in the GSL layer to the actual sequence length for each ECG to facilitate model training in batches. Future work can improve the model implementation to allow using shorter temporal resolutions for biosignals with variable lengths. Lastly, we only investigate biosignal classification tasks in this study. Future study can apply our model to other use cases.

7. Conclusion

In conclusion, we developed GRAPHs4MER, a general GNN for modeling multivariate biosignals. Our method outperformed prior work on seizure detection, sleep staging, and ECG classification, and

learned meaningful graph structures that reflect the underlying classes.

Acknowledgments

We thank Stanford HAI and Google Cloud for the support of Google Cloud credits for this work.

References

- David Ahmedt-Aristizabal, Tharindu Fernando, Simon Denman, Lars Petersson, Matthew J Aburn, and Clinton Fookes. Neural memory networks for seizure type classification. *Conf. Proc. IEEE Eng. Med. Biol. Soc.*, 2020:569–575, July 2020.
- Frédérique Amor, Sylvain Baillet, Vincent Navarro, Claude Adam, Jacques Martinerie, and Michel Le Van Quyen. Cortical local and long-range synchronization interplay in human absence seizure initiation. *Neuroimage*, 45(3):950–962, April 2009.
- Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In Maria Florina Balcan and Kilian Q Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1120–1128, New York, New York, USA, 2016.
- Umar Asif, Subhrajit Roy, Jianbin Tang, and Stefan Harrer. SeizureNet: Multi-Spectral deep feature learning for seizure type classification. In *Machine Learning in Clinical Neuroimaging and Radiogenomics in Neuro-oncology*, pages 77–87. Springer International Publishing, 2020.
- Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. Adaptive graph convolutional recurrent network for traffic forecasting. In *Advances in Neural Information Processing Systems*, volume 33, pages 17804–17815, 2020.
- Belkin and Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. *Adv. Neural Inf. Process. Syst.*, 14, 2001.
- Luc Berthouze, Leon M James, and Simon F Farmer. Human EEG shows long-range temporal correlations of oscillation amplitude in theta, alpha and beta bands across a wide age range. *Clin. Neurophysiol.*, 121(8):1187–1197, August 2010.

- Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Process. Mag.*, 34(4):18–42, July 2017.
- Chen Cai and Yusu Wang. A note on over-smoothing for graph neural networks. *arXiv preprint arXiv:2006.13318*, 2020.
- Ines Chami, Sami Abu-El-Haija, Bryan Perozzi, Christopher Ré, and Kevin Murphy. Machine learning on graphs: A model and comprehensive taxonomy. *Journal of Machine Learning Research*, 23(89):1–64, 2022.
- Yu Chen, Lingfei Wu, and Mohammed Zaki. Iterative deep graph learning for graph neural networks: Better and robust node embeddings. *Adv. Neural Inf. Process. Syst.*, 33:19314–19326, 2020.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar, October 2014.
- Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ua6zuk0WRH>.
- Ian C Covert, Balu Krishnan, Imad Najm, Jiening Zhan, Matthew Shore, John Hixson, and Ming Jack Po. Temporal graph convolutional networks for automatic seizure detection. In Finale Doshi-Velez, Jim Fackler, Ken Jung, David Kale, Rajesh Ranganath, Byron Wallace, and Jenna Wiens, editors, *Proceedings of the 4th Machine Learning for Healthcare Conference*, volume 106 of *Proceedings of Machine Learning Research*, pages 160–180. PMLR, 2019.
- Tri Dao, Daniel Y Fu, Stefano Ermon, Atri Rudra, and Christopher Re. Flashattention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems*, 2022.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems*, 29, 2016.
- Ahmed El-Gazzar, Rajat Mani Thomas, and Guido van Wingen. Dynamic adaptive Spatio-Temporal graph convolution for fMRI modelling. In *Machine Learning in Clinical Neuroimaging*, pages 125–134, 2021.
- Ahmed El-Gazzar, Rajat Mani Thomas, and Guido van Wingen. Improving the diagnosis of psychiatric disorders with self-supervised graph state space models. *arXiv preprint arXiv:2206.03331*, 2022.
- N Benjamin Erichson, Omri Azencot, Alejandro Queiruga, Liam Hodgkinson, and Michael W Mahoney. Lipschitz recurrent neural networks. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=-N7PBXqOUJZ>.
- Jianfei Gao and Bruno Ribeiro. On the equivalence between temporal and static equivariant graph representations. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 7052–7076. PMLR, 2022.
- Xiang Gao, Wei Hu, and Zongming Guo. Exploring Structure-Adaptive graph learning for robust Semi-Supervised classification. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6, July 2020.
- Tomas B Garcia. *12-lead ECG: The Art of Interpretation*. Jones & Bartlett Learning, November 2013.
- P Gloor, M Avoli, and G Kostopoulos. Thalamocortical relationships in generalized epilepsy with bilaterally synchronous spike-and-wave discharge. *Generalized epilepsy: neurobiological approaches*, pages 190–212, 1990.

- Karan Goel, Albert Gu, Chris Donahue, and Christopher Re. It’s raw! Audio generation with state-space models. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 7616–7633. PMLR, 17–23 Jul 2022.
- Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. Hippo: Recurrent memory with optimal polynomial projections. *Advances in neural information processing systems*, 33:1474–1487, 2020.
- Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. Combining recurrent, convolutional, and continuous-time models with linear state space layers. *Advances in neural information processing systems*, 34:572–585, 2021.
- Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022.
- Antoine Guillot and Valentin Thorey. RobustSleepNet: Transfer learning for automated sleep staging at scale. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 29:1441–1451, July 2021.
- Antoine Guillot, Fabien Sauvet, Emmanuel H Durning, and Valentin Thorey. Dreem open datasets: Multi-Scored sleep datasets to compare human and automated sleep staging. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 28(9):1955–1965, September 2020.
- Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 558–567. IEEE, June 2019.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for pre-training graph neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=HJLWWJSFDH>.
- Conrad Iber, Sonia Ancoli-Israel, Andrew L Chesson, and Stuart F Quan. *The AASM manual for the scoring of sleep and associated events: rules, terminology and technical specifications*, volume 1. American academy of sleep medicine, 2007.
- Hassan Ismail Fawaz, Benjamin Lucas, Germain Forestier, Charlotte Pelletier, Daniel F Schmidt, Jonathan Weber, Geoffrey I Webb, Lhassane Idoumghar, Pierre-Alain Muller, and François Petitjean. InceptionTime: Finding AlexNet for time series classification. *Data Min. Knowl. Discov.*, 34(6):1936–1962, November 2020.
- Herbert H Jasper. Ten-twenty electrode system of the international federation. *Electroencephalogr. Clin. Neurophysiol.*, 10:371–375, 1958.
- Ziyu Jia, Youfang Lin, Jing Wang, Ronghao Zhou, Xiaojun Ning, Yuanlai He, and Yaoshuai Zhao. Graphsleepnet: Adaptive spatial-temporal graph convolutional networks for sleep stage classification. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 1324–1330, July 2020.
- Weiwei Jiang and Jiayun Luo. Graph neural network for traffic forecasting: A survey. *Expert Syst. Appl.*, 207:117921, November 2022.
- Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. Graph structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD ’20*, page 66–74, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379984. doi: 10.1145/3394486.3403049. URL <https://doi.org/10.1145/3394486.3403049>.
- Vassilis Kalofolias. How to learn a graph from smooth signals. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 920–929, Cadiz, Spain, 2016. PMLR.
- Amol Kapoor, Xue Ben, Luyang Liu, Bryan Peruzzi, Matt Barnes, Martin Blais, and Shawn O’Banion. Examining covid-19 forecasting using spatio-temporal graph neural networks, 2020. URL <https://arxiv.org/abs/2007.03113>.

- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pages 5156–5165. PMLR, 2020.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *International Conference on Learning Representations*, 2020.
- Jia Li, Zhichao Han, Hong Cheng, Jiao Su, Pengyun Wang, Jianfeng Zhang, and Lujia Pan. Predicting path failure in Time-Evolving graphs. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pages 1279–1289, New York, NY, USA, July 2019. Association for Computing Machinery.
- Ruoyu Li, Sheng Wang, Feiyun Zhu, and Junzhou Huang. Adaptive graph convolutional neural networks. *AAAI*, 32(1), April 2018a.
- Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *International Conference on Learning Representations*, 2018b. URL <https://openreview.net/forum?id=SJiHXGWAZ>.
- Feifei Liu, Chengyu Liu, Lina Zhao, Xiangyu Zhang, Xiaoling Wu, Xiaoyan Xu, Yulin Liu, Caiyun Ma, Shoushui Wei, Zhiqiang He, Jianqing Li, and Eddie Ng Yin Kwee. An open access database for evaluating the algorithms of electrocardiogram rhythm and morphology abnormality detection. *Journal of Medical Imaging and Health Informatics*, 8(7): 1368–1373, 2018.
- Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Skq89Scxx>.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Dongsheng Luo, Wei Cheng, Wenchao Yu, Bo Zong, Jingchao Ni, Haifeng Chen, and Xiang Zhang. Learning to drop: Robust graph neural network via topological denoising. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, WSDM '21, pages 779–787, New York, NY, USA, March 2021. Association for Computing Machinery.
- Christoph M Michel and Micah M Murray. Towards the utilization of EEG as a brain imaging tool. *Neuroimage*, 61(2):371–385, June 2012.
- Antonio Ortega, Pascal Frossard, Jelena Kovačević, José M F Moura, and Pierre Vandergheynst. Graph signal processing: Overview, challenges, and applications. *Proc. IEEE*, 106(5):808–828, May 2018.
- George Panagopoulos, Giannis Nikolentzos, and Michalis Vazirgiannis. Transfer graph neural networks for pandemic forecasting. *AAAI*, 35(6):4838–4845, May 2021.
- Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, and Charles Leiserson. EvolveGCN: Evolving graph convolutional networks for dynamic graphs. *AAAI*, 34(04):5363–5370, April 2020.
- C-K Peng, S Havlin, JM Hausdorff, JE Mietus, HE Stanley, and AL Goldberger. Fractal mechanisms and heart rate dynamics: long-range correlations and their breakdown with disease. *Journal of electrocardiology*, 28:59–65, 1995.
- Huy Phan and Kaare Mikkelsen. Automatic sleep staging of EEG signals: recent development, challenges, and future directions. *Physiol. Meas.*, 43(4), April 2022.
- Ashirbad Pradhan, Jiayuan He, and Ning Jiang. Hand gesture recognition and biometric authentication using a multi-day dataset. In *Intelligent Robotics and Applications*, pages 375–385. Springer International Publishing, 2022.
- Zhilin Qu, Gang Hu, Alan Garfinkel, and James N Weiss. Nonlinear and stochastic dynamics in the heart. *Phys. Rep.*, 543(2):61–162, October 2014.
- Ilene M Rosen, Douglas B Kirsch, Kelly A Carden, Raman K Malhotra, Kannan Ramar, R Nisha Aurora, David A Kristo, Jennifer L Martin, Eric J Olson, Carol L Rosen, James A Rowley, Anita V Shelgikar, and American Academy of Sleep Medicine Board of Directors. Clinical use of a home sleep apnea test: An updated american academy of sleep

- medicine position statement. *J. Clin. Sleep Med.*, 14(12):2075–2077, December 2018.
- Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. Temporal graph networks for deep learning on dynamic graphs. *arXiv preprint arXiv:2006.10637*, 2020.
- Khaled Saab, Jared Dunnmon, Christopher Ré, Daniel Rubin, and Christopher Lee-Messer. Weak supervision as an efficient approach for automated seizure detection in electroencephalography. *NPJ Digit Med*, 3:59, April 2020.
- Mona Sazgar and Michael G Young. Overview of EEG, electrode placement, and montages. In *Absolute Epilepsy and EEG Rotation Review: Essentials for Trainees*, pages 117–125. Springer International Publishing, Cham, 2019.
- Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. Structured sequence modeling with graph convolutional recurrent networks. In *Neural Information Processing: 25th International Conference, ICONIP 2018, Siem Reap, Cambodia, December 13-16, 2018, Proceedings, Part I 25*, pages 362–373. Springer, 2018.
- Vinit Shah, Eva von Weltin, Silvia Lopez, James Riley McHugh, Lillian Veloso, Meysam Golmohammadi, Iyad Obeid, and Joseph Picone. The Temple University Hospital Seizure Detection Corpus. *Front. Neuroinform.*, 12:83, November 2018.
- Chao Shang, Jie Chen, and Jinbo Bi. Discrete graph structure learning for forecasting multiple time series. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=WEHS1H5m0k>.
- Afshin Shoeibi, Marjane Khodatars, Navid Ghassemi, Mahboobeh Jafari, Parisa Moridian, Roohallah Alizadehsani, Maryam Panahiazar, Fahime Khozeimeh, Assef Zare, Hossein Hosseini-Nejad, Abbas Khosravi, Amir F Atiya, Diba Aminshahidi, Sadiq Hussain, Modjtaba Rouhani, Saeid Nahavandi, and Udyavara Rajendra Acharya. Epileptic seizures detection using deep learning techniques: A review. *Int. J. Environ. Res. Public Health*, 18(11), May 2021.
- Nils Strodthoff, Patrick Wagner, Tobias Schaeffter, and Wojciech Samek. Deep learning for ECG analysis: Benchmarks and insights from PTB-XL. *IEEE J Biomed Health Inform.*, 25(5):1519–1528, May 2021.
- Akara Supratak, Hao Dong, Chao Wu, and Yike Guo. DeepSleepNet: A model for automatic sleep stage scoring based on raw Single-Channel EEG. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 25(11):1998–2008, November 2017.
- Jiabin Tang, Tang Qian, Shijing Liu, Shengdong Du, Jie Hu, and Tianrui Li. Spatio-temporal latent graph structure learning for traffic forecasting. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2022a.
- Siyi Tang, Jared Dunnmon, Khaled Kamal Saab, Xuan Zhang, Qianying Huang, Florian Dubost, Daniel Rubin, and Christopher Lee-Messer. Self-Supervised graph neural networks for improved electroencephalographic seizure analysis. In *International Conference on Learning Representations*, April 2022b.
- Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena: A benchmark for efficient transformers, 2020. URL <https://arxiv.org/abs/2011.04006>.
- Chenyu Tian and Wai Kin (victor) Chan. Spatial-temporal attention wavenet: A deep learning framework for traffic prediction considering spatial-temporal dependencies. *IET Intel. Transport Syst.*, 15(4):549–561, April 2021.
- Vaswani, Shazeer, Parmar, and others. Attention is all you need. *Adv. Neural Inf. Process. Syst.*, 2017.
- Zhiguang Wang, Weizhong Yan, and Tim Oates. Time series classification from scratch with deep neural networks: A strong baseline. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 1578–1585, May 2017.
- Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep spatial-temporal graph modeling. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 1907–1913. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/264. URL <https://doi.org/10.24963/ijcai.2019/264>.

Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. Inductive representation learning on temporal graphs. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rJeW1yHYwH>.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ryGs6iA5Km>.

Qi Zhang, Jianlong Chang, Gaofeng Meng, Shiming Xiang, and Chunhong Pan. Spatio-Temporal graph structure learning for traffic forecasting. *AAAI*, 34(01):1177–1185, April 2020.

Xiang Zhang and Marinka Zitnik. GNNGuard: Defending graph neural networks against adversarial attacks. *Adv. Neural Inf. Process. Syst.*, 33:9263–9275, 2020.

Xiang Zhang, Marko Zeman, Theodoros Tsiligkaridis, and Marinka Zitnik. Graph-guided network for irregularly sampled multivariate time series. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=Kwm8I7dU-15>.

Cheng Zheng, Bo Zong, Wei Cheng, Dongjin Song, Jingchao Ni, Wenchao Yu, Haifeng Chen, and Wei Wang. Robust graph representation learning via neural sparsification. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11458–11468. PMLR, 2020a.

Chuanpan Zheng, Xiaoliang Fan, Cheng Wang, and Jianzhong Qi. GMAN: A graph Multi-Attention network for traffic prediction. *AAAI*, 34(01):1234–1241, April 2020b.

Yanqiao Zhu, Weizhi Xu, Jinghao Zhang, Yuanqi Du, Jiayu Zhang, Qiang Liu, Carl Yang, and Shu Wu. A survey on graph structure learning: Progress and opportunities. *arXiv e-prints*, pages arXiv–2103, 2021.

Appendix A. Pseudocode for GraphS4mer

Algorithm 1 shows the pseudocode for GRAPHs4MER. Note that $n_d = \frac{T}{r}$ is the number of dynamic graphs, where T is the sequence length and r is the pre-specified resolution. While there is a for loop over $t = 1, \dots, n_d$, the code implementation can be done in batches to eliminate the need of a for loop.

Appendix B. Details of datasets and baselines

Table 5 summarizes the task, sequence length, and number of nodes in each of the three datasets.

Temple University Hospital Seizure Detection Corpus (TUSZ). We use the publicly available Temporal University Hospital Seizure Detection Corpus (TUSZ) v1.5.2. for seizure detection (Shah et al., 2018). We follow the same experimental setup as in Tang et al. (2022b). The TUSZ train set is divided into train and validation splits with distinct patients, and the TUSZ test set is held-out for model evaluation. The following 19 EEG channels are included: FP1, FP2, F3, F4, C3, C4, P3, P4, O1, O2, F7, F8, T3, T4, T5, T6, FZ, CZ, and PZ. Because the EEG signals are sampled at different sampling rate, we re-sample all the EEG signals to 200 Hz. Following a prior study (Tang et al., 2022b), we also exclude 5 patients in the TUSZ test set who appear in both the TUSZ train and test sets. The EEG signals are divided into 60-s EEG clips without overlaps, and the task is to predict whether or not an EEG clip contains seizure. Table 6 shows the number of EEG clips and patients in the train, validation, and test splits.

We compare our model performance to existing models for seizure detection, including (1) LSTM (Hochreiter and Schmidhuber, 1997), a variant of RNN with gating mechanisms; (2) Dense-CNN (Saab et al., 2020), a densely connected CNN specialized in seizure detection; (3) CNN-LSTM (Ahmedt-Aristizabal et al., 2020); and (4) Dist- and Corr-DCRNN without and with self-supervised pre-training (Tang et al., 2022b).

Dreem Open Dataset-Healthy (DOD-H). We use the publicly available Dreem Open Dataset-Healthy (DOD-H) for sleep staging (Guillot et al., 2020). The DOD-H dataset consists of overnight PSG sleep recordings from 25 volunteers. The PSG signals are measured from 12 EEG channels, 1 elec-

Algorithm 1: Pseudocode for GRAPHs4MER.

Input : multivariate biosignal \mathbf{X} , label \mathbf{y}
Parameter: model weights θ , including weights for S4, GSL, GNN, and linear layers
Output : prediction $\hat{\mathbf{y}}$, learned graph adjacency matrices $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(n_d)}$
 Randomly initialize model weights θ
while not converged do
 $\mathcal{L}_{\text{reg}} \leftarrow 0$
 // step 1. get S4 embeddings
 $\mathbf{H} \leftarrow \text{S4}(\mathbf{X})$
 for $t = 1, 2, \dots, n_d$ **do**
 $\mathbf{h}^{(t)} \leftarrow \text{mean-pool}(\mathbf{H}_{:,((t-1) \times r):(t \times r),:})$
 // mean-pool along temporal dim
 // step 2. graph structure learning (GSL)
 Construct KNN graph, $\mathbf{W}_{\text{KNN}}^{(t)}$, based on cosine similarity between nodes using $\mathbf{h}^{(t)}$
 $\mathbf{W}^{(t)} \leftarrow \text{GSL}(\mathbf{h}^{(t)}, \mathbf{W}_{\text{KNN}}^{(t)})$ using Eqn. 7
 -9
 $\mathcal{L}_{\text{reg}} \leftarrow \mathcal{L}_{\text{reg}} + \alpha \mathcal{L}_{\text{smooth}}(\mathbf{h}^{(t)}, \mathbf{W}^{(t)}) + \beta \mathcal{L}_{\text{degree}}(\mathbf{W}^{(t)}) + \gamma \mathcal{L}_{\text{sparse}}(\mathbf{W}^{(t)})$
 using Eqn. 4-6 // graph regularization loss
 // step 3. get node embeddings from GNN
 $\mathbf{z}^{(t)} \leftarrow \text{GNN}(\mathbf{h}^{(t)}, \mathbf{W}^{(t)})$ // $\mathbf{z}^{(t)}$ has shape $(N, 1, D_{\text{hidden}})$
 end
 $\mathbf{Z} \leftarrow \text{concat}(\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \dots, \mathbf{z}^{(n_d)})$
 // concatenate along temporal dim
 // step 4. pooling and fully connected layers
 $\mathbf{Z} \leftarrow \text{graph-pool}(\text{temporal-pool}(\mathbf{Z}))$
 $\hat{\mathbf{y}} \leftarrow \text{Linear}(\mathbf{Z})$
 $\mathcal{L}_{\text{pred}} \leftarrow \mathcal{L}_{\text{cross-entropy}}(\hat{\mathbf{y}}, \mathbf{y})$ // loss for classification
 $\mathcal{L} \leftarrow \mathcal{L}_{\text{pred}} + \frac{1}{n_d} \mathcal{L}_{\text{reg}}$ // total loss
 Back-propagate \mathcal{L} to update model weights θ
end

tromyographic (EMG) channel, 2 electrooculography (EOG) channels, and 1 electrocardiogram channel using a Siesta PSG device (Compumedics). All the signals are sampled at 250 Hz. Following the standard AASM Scoring Manual and Recommendations (Rosen et al., 2018), each 30-s PSG signal is annotated by a consensus of 5 experienced sleep technologists as one of the five sleep stages: wake, rapid eye movement (REM), non-REM sleep stages N1, N2, and N3. We randomly split the PSG signals by 60/20/20 into train/validation/test splits, where each split has distinct subjects. Table 7 shows the number of 30-s PSG clips, the number of subjects, and the five sleep stage distributions.

Baseline models include existing sleep staging models that achieved state-of-the-art on DOD-H dataset, SimpleSleepNet (Guillot et al., 2020), RobustSleepNet (Guillot and Thorey, 2021), and DeepSleepNet (Supratak et al., 2017), all of which are based on CNNs and/or RNNs. We also include the traditional sequence model LSTM (Hochreiter and Schmidhuber, 1997) as a baseline. For fair comparisons between the baselines and GRAPHs4MER, we trained SimpleSleepNet, RobustSleepNet, and DeepSleepNet using the open sourced code¹ and set the temporal context to be one 30-s PSG signal.

ICBEB ECG Dataset. We use the publicly available ECG dataset for the 1st China Physiological Signal Challenge held during the International Conference on Biomedical Engineering and Biotechnology (ICBEB) for ECG classification (Liu et al., 2018). We follow the same data split as described in Strodtzoff et al. (2021). Specifically, as the official ICBEB test set is not publicly available, we only use the official ICBEB training set and randomly split it into 10 folds using stratified split, where the first 8 folds are used as the training set, the 9th fold is used as the validation set for hyperparameter tuning, and the 10th fold is used as the held-out test set to report results in this study. In total, there are 6,877 12-lead ECGs ranging between 6-s and 60-s. Following Strodtzoff et al. (2021), the ECGs are downsampled to 100 Hz, resulting in sequence lengths ranging from 600 to 6k time steps. To handle variable length ECGs for training in mini batches, we pad the short sequences with 0s. During training and testing, the padded values are masked out so that they are not seen by the models.

1. SimpleSleepNet: <https://github.com/Dreem-Organization/dreem-learning-open>; RobustSleepNet, DeepSleepNet: <https://github.com/Dreem-Organization/RobustSleepNet>

Table 5: Summary of datasets.

Dataset	Biosignal Type	Task	Sequence Length (#Time steps)	Number of Nodes
TUSZ	EEG	Binary classification of seizure vs non-seizure	12,000	19
DOD-H	PSG	Multiclass classification of five sleep stages	7,500	16
ICBEB	ECG	Multilabel classification of nine ECG classes	Varying from 600 to 6,000	12

Table 6: Number of EEG clips and patients in train, validation, and test splits of TUSZ dataset.

	EEG Clips (% Seizure)	Patients (% Seizure)
Train Set	38,613 (9.3%)	530 (34.0%)
Validation Set	5,503 (11.4%)	61 (36.1%)
Test Set	8,848 (14.7%)	45 (77.8%)

Table 7: Number of subjects and 30-s PSG clips in the train, validation, and test splits of DOD-H dataset.

	Subjects	30-s PSG Clips					
		Total	Wake	N1	N2	N3	REM
Train	15	14,823	1,839	925	6,965	2,015	3,079
Validation	5	5,114	480	254	2,480	990	910
Test	5	4,725	718	326	2,434	509	738

In the ICBEB dataset, each ECG record is annotated by up to three reviewers. There are 9 classes in total, including one normal and 8 abnormal classes, and each ECG may be associated with more than one abnormal classes (i.e., multilabel classification). Table 8 shows the 9 classes and the number of ECGs in each class in the train/validation/test splits.

We compare GRAPHs4MER to a wide variety of prior CNNs/RNNs for ECG classification as in [Strodthoff et al. \(2021\)](#): (1) InceptionTime ([Ismail Fawaz et al., 2020](#)), (2) XResNet1D ([He et al., 2019](#)), (3) ResNet1D ([Wang et al., 2017](#)), (4) fully convolutional network (FCN) ([Wang et al., 2017](#)), (5) LSTM ([Hochreiter and Schmidhuber, 1997](#)), (6) bidirectional LSTM ([Hochreiter and Schmidhuber, 1997](#)), and (7) WaveletNN ([Strodthoff et al., 2021](#)). Note that these baselines take 2.5-s ECG windows as inputs and aggregate the window-wise predictions at test time, whereas GRAPHs4MER takes the entire ECG signal as input, which allows modeling long-range temporal dependencies in the entire signals. We ran the baseline models using the open source im-

plementations² for three runs with different random seeds.

Table 8: Number of ECGs in the train, validation, and test splits of ICBEB dataset used in this study. Note that some ECGs are associated with more than one abnormal classes, and thus the total number of ECGs (last row) is not equal to the sum of the ECGs in the individual classes. *Abbreviations*: AFIB, atrial fibrillation; I-AVB, first-degree atrioventricular block; LBBB, left bundle branch block; RBBB, right bundle branch block; PAC, premature atrial contraction; PVC, premature ventricular contraction; STD, ST-segment depression; STE, ST-segment elevated.

	Train	Validation	Test
Normal	734	92	92
AFIB	976	122	123
I-AVB	578	72	72
LBBB	189	24	23
RBBB	1,487	176	194
PAC	493	61	62
PVC	560	70	70
STD	695	87	87
STE	176	22	22
Total	5,499	690	688

Appendix C. Details of model training procedures and hyperparameters

Model training was accomplished using the AdamW optimizer ([Loshchilov and Hutter, 2019](#)) in PyTorch on a single NVIDIA A100 GPU. All experiments were

2. https://github.com/helme/ecg_ptbxx1_benchmarking

run for three runs with different random seeds. Cosine learning rate scheduler with 5-epoch warm start was used (Loshchilov and Hutter, 2017). Model training was early stopped when the validation loss did not decrease for 20 consecutive epochs. We performed the following hyperparameter search on the validation set using an off-the-shelf hyperparameter tuning tool³: (1) initial learning rate within range [1e-4, 1e-2]; (2) dropout rate in S4 and GNN layers within range [0.1, 0.5]; (3) hidden dimension of S4 and GNN layers within range {64, 128, 256}; (4) number of S4 layers within range {2, 3, 4}; (5) S4 bidirectionality; (6) number of GNN layers within range {1, 2}; (7) graph pooling within range {mean-pool, max-pool, sum-pool}; (8) value of κ threshold for graph pruning by keeping edges whose weights $> \kappa$ (κ within range [0.01, 0.5]); (9) KNN graph with $K \in \{2, 3\}$; (10) weight of KNN graph $\epsilon \in [0.3, 0.6]$; (11) α , β , and γ weights in graph regularization within range [0, 1].

Model training and hyperparameters for seizure detection on TUSZ dataset. As there are many more negative samples in the dataset, we undersampled the negative examples in the train set during training. We used binary cross-entropy loss as the loss function. The models were trained for 100 epochs with an initial learning rate of 8e-4. The batch size was 4; dropout rate was 0.1; hidden dimension was 128; number of stacked S4 layers was 4; S4 layers were unidirectional; number of GNN layers was 1; graph pooling was max-pool and temporal pooling was mean-pool; graph pruning was done by setting a threshold of $\kappa = 0.1$, where edges whose edge weights ≤ 0.1 were removed; $K = 2$ for KNN graph and the KNN graph weight ϵ was 0.6; α , β , and γ weights were all set to 0.05. This results in 265k trainable parameters in GRAPHS4MER. Best model was picked based on the highest AUROC on the validation set. To obtain binarized predictions, we selected the cutoff threshold by maximizing the F1-score on the validation set.

Model training and hyperparameters for sleep staging on DOD-H dataset. As the DOD-H dataset is highly imbalanced, we undersampled the majority classes in the train set during training. We used cross-entropy loss as the loss function. The models were trained for 100 epochs with an initial learning rate of 1e-3. The batch size was 4; dropout rate was 0.4; hidden dimension was 128; number of stacked S4

layers was 4 and S4 layers were unidirectional; number of GNN layers was 1; graph pooling was sum-pool and temporal pooling was mean-pool; graph pruning was done by setting a threshold of $\kappa = 0.1$; $K = 3$ for KNN graph and the weight for KNN graph ϵ was 0.6; α , β , and γ weights were all set to 0.2. This results in 266k trainable parameters in GRAPHS4MER. Best model was picked based on the highest macro-F1 score on the validation set.

Model training and hyperparameters for ECG classification on ICBEb dataset. As each ECG in the ICBEb dataset may be associated with more than one classes, this task is a multilabel classification task. Therefore, we used binary cross-entropy loss as the loss function. The models were trained for 100 epochs with an initial learning rate of 1e-3. The batch size was 8; dropout rate was 0.1; hidden dimension was 128; number of stacked S4 layers was 4; S4 layers were bidirectional; number of GNN layers was 1; graph pooling was mean-pool and temporal pooling was mean-pool; graph pruning was done by setting a threshold of $\kappa = 0.02$; $K = 2$ for KNN graph and the weight for KNN graph ϵ was 0.6; α , β , and γ weights were 1.0, 0.0, and 0.5, respectively. This results in 299k trainable parameters in GRAPHS4MER for ECG classification. Best model was picked based on the highest macro-AUROC on the validation set. To obtain binarized predictions, we selected the cutoff threshold for each class separately by maximizing the F1-score on the validation set for the respective class.

Appendix D. Evaluation metrics for ECG classification on ICBEb dataset

In binary classification, F_β score uses a positive real factor β to weigh recall β times as important as precision:

$$F_\beta = (1 + \beta^2) \times \frac{\text{Precision} \times \text{Recall}}{(\beta^2 \times \text{Precision}) + \text{Recall}}$$

Specifically, F1-score (i.e., $\beta = 1$) is the harmonic mean of precision and recall:

$$F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

F2-score considers recall more important than precision:

$$F2 = \frac{(1 + 2^2) \times \text{Precision} \times \text{Recall}}{(2^2 \times \text{Precision}) + \text{Recall}}$$

3. <https://docs.wandb.ai/guides/sweeps>

In contrast, G_β measure is defined as:

$$G_\beta = \frac{TP}{(TP + FP + \beta \times FN)}$$

where TP, FP, and FN are the number of true positives, false positives, and false negatives, respectively. G_2 -score is calculated by setting $\beta = 2$ in the above equation.

For multilabel/multiclass classification, macro-averaged scores are computed by taking the mean of the individual scores in the individual classes. For ECG classification experiments, we report macro-F1, macro-F2, macro-G2, and macro-AUROC in Table 3.

Appendix E. Model performance on individual ECG classes

Table 9 shows GRAPHs4MER and baseline model performance for individual ECG classes. GRAPHs4MER outperforms the baselines on 6 out of 9 classes.

Appendix F. Additional evaluation metrics for ablation experiments

Table 10 shows additional evaluation metrics for ablation experiments on TUSZ, DOD-H, and ICBEB datasets.

Appendix G. Effect of temporal resolution

To examine the effect of temporal resolution r on model performance, we show the performance versus different values of r for seizure detection and sleep staging in Figure 3. We observe that smaller value of r tends to result in higher performance, which suggests that capturing dynamically varying graph structures is useful for these tasks.

Appendix H. Effect of KNN graph weight and graph regularization

Figure 4 shows the validation set performance for GRAPHs4MER with respect to different values of KNN graph weight (i.e., ϵ in Equation 9). A larger KNN graph weight results in marginally better model performance.




Figure 3: Model performance versus temporal resolution using the run with median performance. For convenience, we assume that temporal resolution is chosen so that the sequence length is divisible by the resolution. Asterisk indicates the temporal resolution used to report results for GRAPHs4MER in Tables 1–2 and Table 4. Error bars indicate 95% confidence intervals obtained using bootstrapping with 5,000 replicates with replacement.

Figure 5 shows the validation set performance for GRAPHs4MER without and with graph regularization loss. Graph regularization results in marginally better model performance.




Figure 4: Model performance on the validation set versus KNN graph weight (ϵ in Equation 9) using the run with median performance. Error bars indicate 95% confidence intervals obtained using bootstrapping with 5,000 replicates with replacement.

Appendix I. Computational cost of GSL

In Table 11, we compare the additional number of parameters and multiply-accumulate operations (MACs) of GRAPHs4MER with respect

Table 9: Model performance (AUROC) for individual ECG classes. Best and second best results are **bolded** and underlined, respectively.

Model	Normal	AFIB	I-AVB	LBBS	RBBB	PAC	PVC	STD	STE
InceptionTime	<u>0.962 ± 0.001</u>	0.973 ± 0.004	0.992 ± 0.001	0.972 ± 0.018	<u>0.990 ± 0.003</u>	0.913 ± 0.022	0.970 ± 0.004	<u>0.956 ± 0.004</u>	0.946 ± 0.025
XResNet1D	0.962 ± 0.000	0.977 ± 0.001	0.985 ± 0.002	<u>0.973 ± 0.004</u>	0.991 ± 0.001	<u>0.955 ± 0.003</u>	<u>0.978 ± 0.002</u>	0.948 ± 0.003	0.966 ± 0.012
ResNet1D	0.960 ± 0.003	<u>0.976 ± 0.002</u>	<u>0.993 ± 0.002</u>	0.963 ± 0.019	0.989 ± 0.001	0.905 ± 0.007	0.975 ± 0.004	0.951 ± 0.005	<u>0.957 ± 0.014</u>
FCN	0.955 ± 0.002	0.972 ± 0.002	0.986 ± 0.002	0.956 ± 0.008	<u>0.990 ± 0.000</u>	0.834 ± 0.013	0.971 ± 0.002	0.953 ± 0.000	0.929 ± 0.013
LSTM	0.903 ± 0.007	0.973 ± 0.004	0.984 ± 0.003	0.934 ± 0.009	0.976 ± 0.004	0.886 ± 0.014	0.942 ± 0.008	0.954 ± 0.007	0.929 ± 0.040
Bidir-LSTM	0.917 ± 0.003	0.975 ± 0.005	0.991 ± 0.001	0.937 ± 0.012	0.984 ± 0.003	0.899 ± 0.009	0.945 ± 0.002	0.946 ± 0.003	0.915 ± 0.035
WaveletNN	0.945 ± 0.002	0.921 ± 0.006	0.901 ± 0.002	0.965 ± 0.005	0.947 ± 0.002	0.737 ± 0.005	0.928 ± 0.004	0.922 ± 0.003	0.932 ± 0.015
GRAPHS4MER (ours)	0.989 ± 0.003	0.970 ± 0.002	0.996 ± 0.001	0.978 ± 0.006	0.986 ± 0.001	0.965 ± 0.004	0.983 ± 0.002	0.971 ± 0.005	0.952 ± 0.007

Table 10: Additional evaluation metrics for ablation experiments. Best and second best results are **bolded** and underlined, respectively.

Model	F1-Score	TUSZ			DOD-H Kappa	ICBEB		
		AUPRC	Sensitivity	Specificity		Macro F2-Score	Macro G2-Score	Macro AUROC
S4	0.491 ± 0.033	0.488 ± 0.017	0.441 ± 0.077	0.941 ± 0.020	0.757 ± 0.012	0.789 ± 0.004	0.577 ± 0.006	0.966 ± 0.004
GRAPHS4MER w/o S4	0.373 ± 0.100	0.302 ± 0.117	0.537 ± 0.133	0.765 ± 0.065	0.597 ± 0.064	0.313 ± 0.006	0.107 ± 0.002	0.529 ± 0.028
GSL-GNN-S4	0.637 ± 0.010	0.668 ± 0.019	0.629 ± 0.053	0.941 ± 0.014	<u>0.784 ± 0.014</u>	0.782 ± 0.016	0.564 ± 0.019	0.966 ± 0.001
GRAPHS4MER w/o GSL	<u>0.672 ± 0.029</u>	<u>0.672 ± 0.029</u>	<u>0.708 ± 0.020</u>	0.931 ± 0.019	0.734 ± 0.010	<u>0.794 ± 0.011</u>	<u>0.588 ± 0.015</u>	0.970 ± 0.007
GRAPHS4MER (ours)	0.680 ± 0.012	0.723 ± 0.023	0.718 ± 0.041	<u>0.933 ± 0.009</u>	0.802 ± 0.014	0.804 ± 0.004	0.609 ± 0.005	0.977 ± 0.001

Table 11: Number of additional parameters and multiply-accumulate operations (MACs) with respect to GRAPHS4MER w/o GSL (with predefined graphs). The additional parameters and MACs come from the self-attention layer in GSL, and the MACs increase as the time interval r decreases. r indicates the time interval (resolution) in the GSL layer, T indicates the sequence length (dependent on dataset), and n_d indicates the number of dynamic graphs. For convenience, we assume that the resolution r is chosen so that the sequence length T is divisible by r . “-” indicates that the sequence length is not divisible by r . “NA” indicates not applicable, as we set r to be the actual sequence lengths of ECGs in the ICBEB dataset due to variable sequence lengths. Asterisk indicates the resolution used to report results for GRAPHS4MER in Tables 1-4.

Model	TUSZ ($T = 12,000$)		DOD-H ($T = 7,500$)		ICBEB (varying T)	
	#Parameters	MACs	#Parameters	MACs	#Parameters	MACs
GRAPHS4MER w/o GSL	ref	ref	ref	ref	ref	ref
GRAPHS4MER ($r = T, n_d = 1$)	32.768K	19.923M	32.768K	16.777M	32.768K*	12.583M*
GRAPHS4MER ($r = T/2, n_d = 2$)	32.768K	39.846M	32.768K	33.554M	NA	NA
GRAPHS4MER ($r = T/3, n_d = 3$)	32.768K	59.769M	32.768K*	50.332M*	NA	NA
GRAPHS4MER ($r = T/4, n_d = 4$)	32.768K	79.692M	32.768K	67.109M	NA	NA
GRAPHS4MER ($r = T/5, n_d = 5$)	32.768K	99.615M	32.768K	83.886M	NA	NA
GRAPHS4MER ($r = T/6, n_d = 6$)	32.768K*	119.538M*	32.768K	100.663M	NA	NA
GRAPHS4MER ($r = T/8, n_d = 8$)	32.768K	159.384M	-	-	NA	NA
GRAPHS4MER ($r = T/10, n_d = 10$)	32.768K	199.229M	32.768K	167.772M	NA	NA

to GRAPHS4MER without GSL (with predefined graphs), as well as computational costs of GRAPHS4MER with different values of resolutions r in the GSL layer (see Section 4.3). The GSL layer (i.e., self-attention) adds additional 33k parameters (approximately 12% of total parameters). Moreover, as the time interval r decreases (i.e., number of

dynamic graphs increases), the number of MACs increases, whereas the number of parameters remains the same. This is expected given that the same GSL layer is applied multiple times to obtain dynamic graphs. Specifically, the additional MACs are proportional to the number of dynamic graphs in the GSL layer.




Figure 5: Model performance on the validation set for GRAPH\$4MER without and with graph regularization loss using the run with median performance. Error bars indicate 95% confidence intervals obtained using bootstrapping with 5,000 replicates with replacement.

Appendix J. Visualization of PSG graphs for five sleep stages

Figure 6 shows mean adjacency matrices for PSG signals in correctly predicted test samples grouped by five sleep stages. We observe that N3 (Figure 6ix) differs from wake stage more than N1 (Figure 6vii) (also see Figure 2). Moreover, in REM stage (Figure 6ii), the EMG channel has very weak connection to all other channels (red arrow in Figure 6ii), which is expected given that one experiences muscle paralysis in REM stage.

Appendix K. Visualization of example ECG signals and ECG graphs

Figure 7 shows example normal ECG, ECG with first-degree atrioventricular block (I-AVB), and ECG with left bundle branch block (LBBB).

Figure 8 shows mean adjacency matrices for ECG signals in correctly predicted test samples, grouped by nine ECG classes. δ indicates the mean and standard deviation of the absolute values of the differences between the abnormal class and the normal class. We observe that the magnitude of difference between I-AVB and normal ECG is the smallest (i.e., smallest δ), whereas the magnitude of difference between LBBB and normal ECG is the largest (i.e., largest δ).




Figure 6: (i)–(v) Mean adjacency matrices for PSG signals for five sleep stages in correctly predicted test samples. (vi)–(ix) Difference between non-WAKE stages and WAKE. δ indicates the mean and standard deviation of the absolute values of the differences between non-WAKE stages and WAKE. Red arrow indicates EMG channel in REM stage that has weak connection to all other channels. Self-edges (i.e., diagonal) are not shown here.

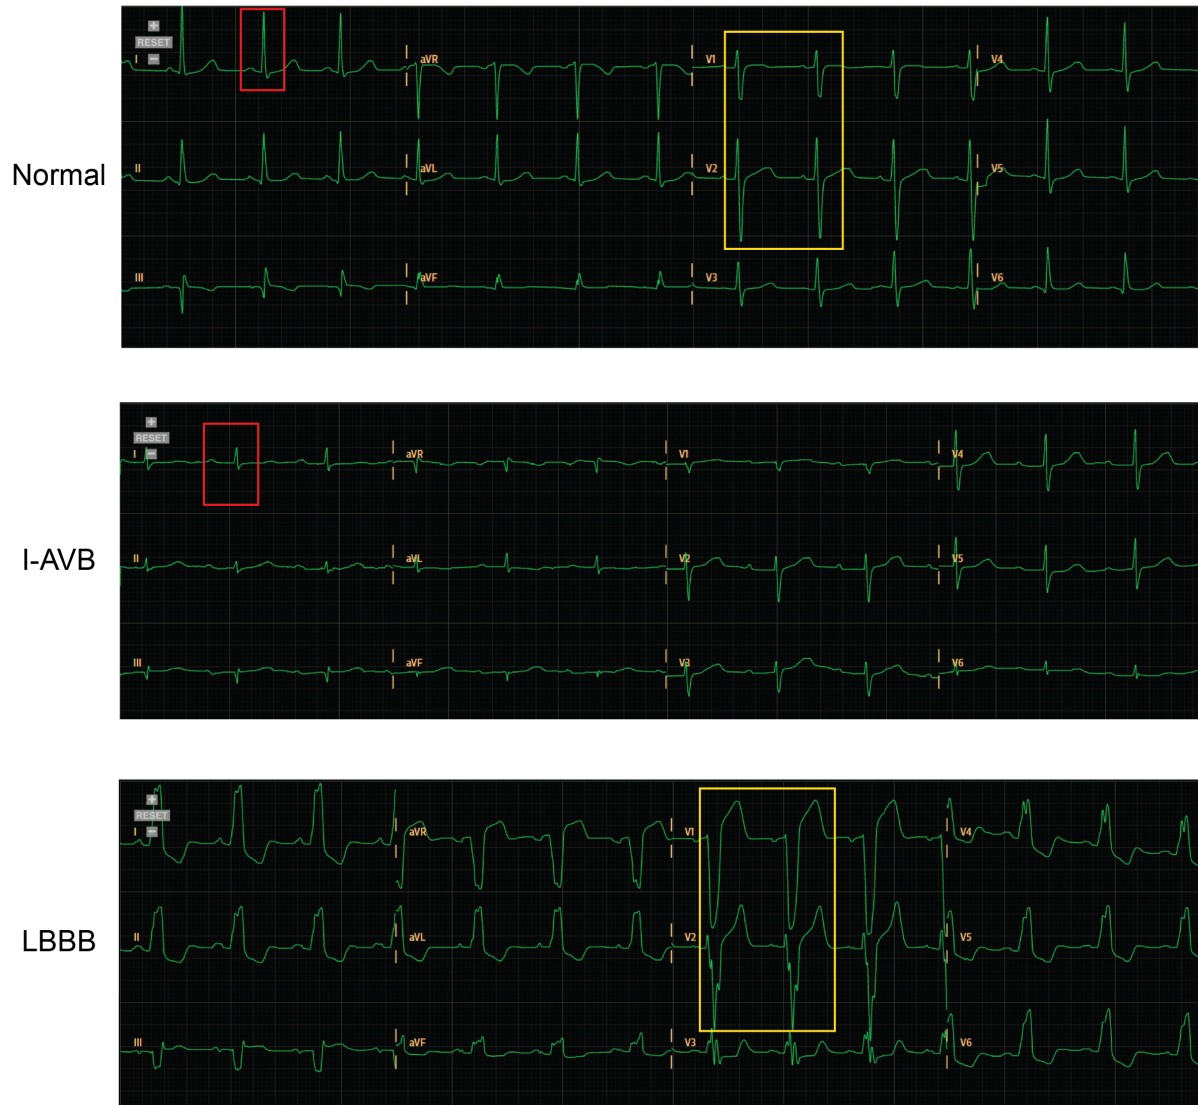


Figure 7: Example normal ECG, ECG with first-degree atrioventricular block (I-AVB), and ECG with left bundle branch block (LBBB). Red boxes demonstrate the difference in PR interval between I-AVB and normal ECG. Yellow boxes demonstrate the pronounced morphology difference in the QRS complex between LBBB and normal ECG.

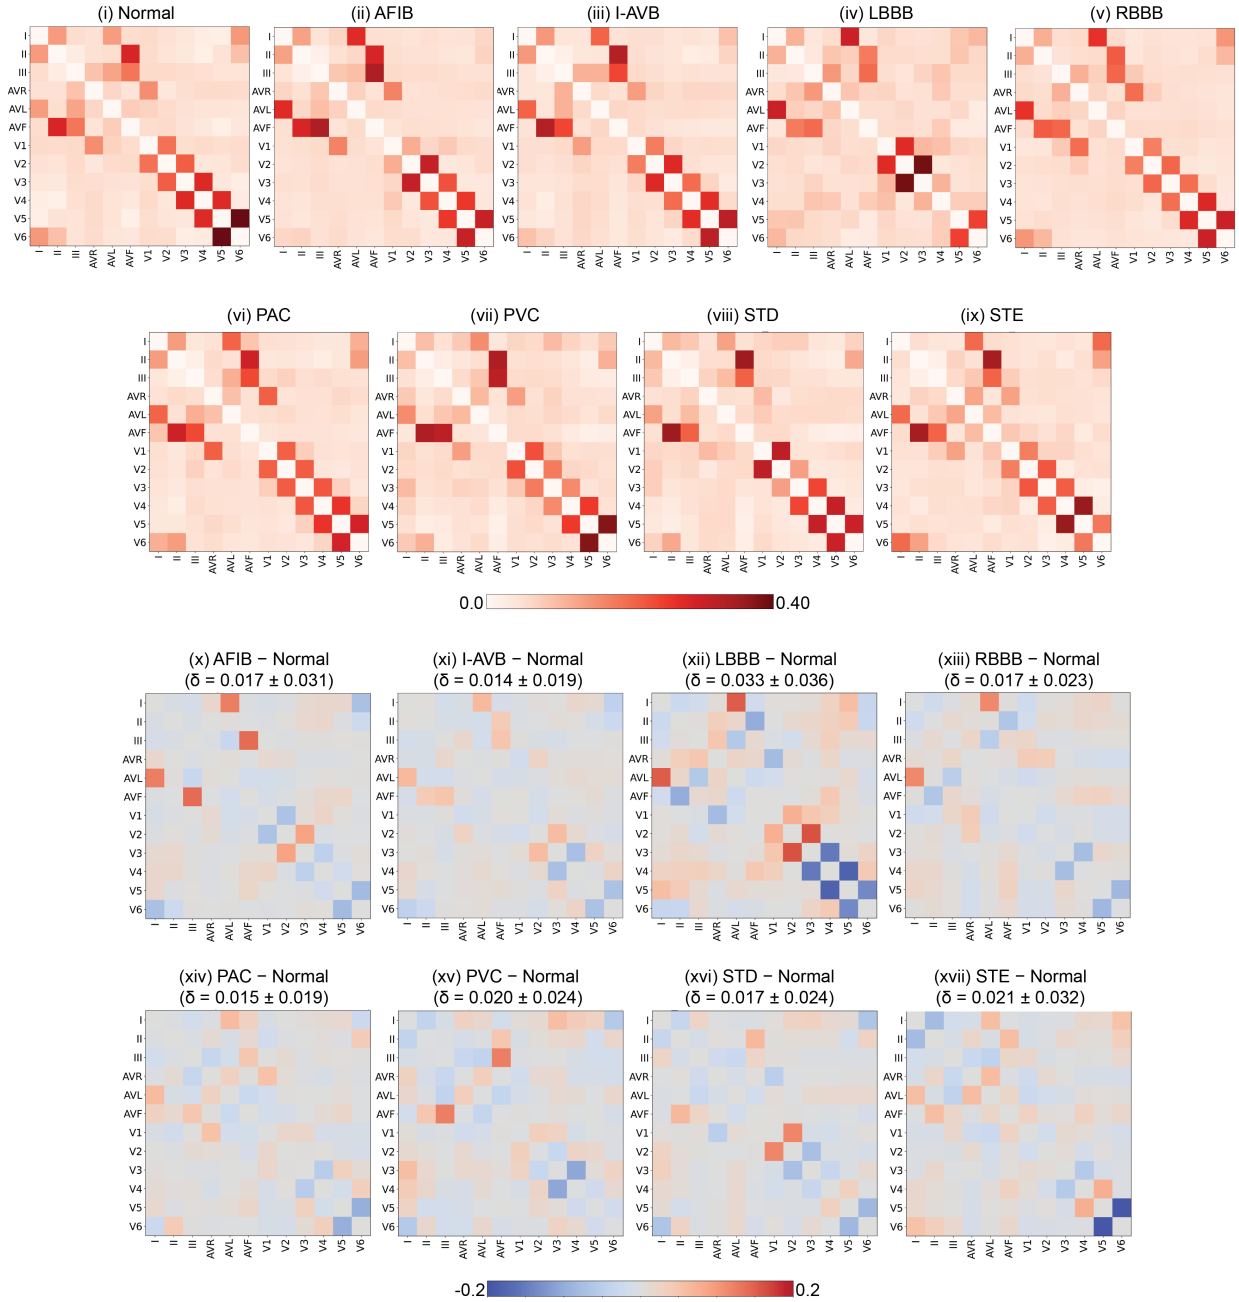


Figure 8: (i)–(ix) Mean adjacency matrices for ECG signals for nine ECG classes in correctly predicted test samples. (x)–(xvii) Difference between adjacency matrices of abnormal ECG classes and normal class. δ indicates the mean and standard deviation of the absolute values of the differences between the abnormal class and the normal class. Self-edges (i.e., diagonal) are not shown here. *Abbreviations:* AFIB, atrial fibrillation; I-AVB, first-degree atrioventricular block; LBBB, left bundle branch block; RBBB, right bundle branch block; PAC, premature atrial contraction; PVC, premature ventricular contraction; STD, ST-segment depression; STE, ST-segment elevated.