LEARNING SELF-SUPERVISED DYNAMIC NETWORKS FOR SEIZURE ANALYSIS

William Cappelletti
LTS4, EPFL
Lausanne, Switzerland

Yuhan Xie
LTS4, EPFL
Lausanne, Switzerland

Pascal Frossard
LTS4, EPFL
Lausanne, Switzerland

ABSTRACT

Recurrent Graph Neural Networks are very effective at modeling brain activity, thanks to their spatial-temporal inductive bias, and they show further capabilities when we apply self-supervised pretraining methods. For instance, they show improved performances on epileptic-seizure analysis, namely detection and classification, compared to convolutional and classical recurrent neural networks. Still, the graphs used by these methods are generally predefined, and often provide little insight on the task. We build upon current advancements in graph learning for time series forecasting to propose a novel architecture to learn task-specific networks, jointly with a self-supervised pretraining strategy. We study the performances of learned graphs at different scale, by comparing static and dynamic networks, and illustrate the outstanding performance of our model on epilepsy classification and detection tasks.

1 INTRODUCTION

Epilepsy affects nearly 50 million people worldwide, according to a 2019 World Health Organization report (World Health Organization, 2019). It is a cerebral disease consisting of recurrent *seizures*, i.e. transient occurrence of signs and/or symptoms due to abnormal excessive or synchronous neuronal activity in the brain Fisher et al. (2005). It is estimated that only 70% of the people with epilepsy could live "seizure-free" if properly diagnosed and treated, either through medicine assumption or brain surgery (World Health Organization, 2019). Being able to reliably detect, and if possible predict, their attacks could improve the life standard of the untreatable subjects.

The main non-intrusive tool used by doctors to monitor the neuronal activity of people with epilepsy is Electroencephalography Schomer & da Silva (2012), which captures neuronal activity with electrodes placed over the scalp. More precisely, we measure voltage fluctuations resulting from ionic current within neurons, which provide us multiple time series, or *channels*, which together constitute an *electroencephalogram* (EEG). Neurologists observe these multivariate signals on a screen, looking for epileptic patterns and manually annotate *onset* and *termination* times of seizures. Furthermore, there are different kinds of seizures, with distinct symptoms, and different patterns and propagation dynamics.

Reliable and explainable machine learning methods would be greatly beneficial for both doctors and patients. Automated seizure analysis consists of detection and classification tasks. The objective of *seizure detection* is to provide a binary label identifying seizures from regular brain activity (or background), while *seizure classification* is a multiclass problem. Still, seizure events are very sparse over EEG sessions, and the different kinds are generally imbalanced. Therefore, it is critical to efficiently leverage the large available amount of background signals.

Deep-learning-based approaches are increasingly used in EEG analysis (Roy et al., 2019), and recent works show good performances on seizure-related tasks (Nhu et al., 2022; Lee et al., 2022). Some of the most successful architectures are convolutional (CNN) and recurrent neural networks (RNNs), which are designed to work with sequential data. Nonetheless, these methods consider each time series as independent, and do not leverage the spatial information that is inherent to EEGs, not the functional connectivity of the brain. Graphs, which are structures that define pairwise relations between entities, provide a natural representation for the structural or functional dependencies between measured brain areas, and Graph Neural Networks (GNNs) have shown impressive performances on seizure analysis.

In this study, we focus on the public Temple University Hospital EEG Seizure Corpus (TUSZ) v1.5.2 (Obeid & Picone, 2016). This dataset consists of over 50000 minutes of EEG signals from 637 patients split in predefined training and test sets. It contains 3050 annotated seizures which span 6.3% of the training and 9.8% of the test data. We divide seizures in four categories, namely combined focal (CF), generalized nonspecific (GN), absence (AB), and combined tonic (CT), which amount respectively to 2165, 523, 99, and 109 events of various durations.

Tang et al. (2021) show the potential of Recurrent Graph Neural Networks for analyzing epileptic seizures over consecutive windows of EEGs signals. In particular, the latent graph representations learned by a DCRNN model (Li et al., 2018) greatly improves detection and classification performances compared to CNNs and LTSMs, which ignore network structure. Still, their main contribution lies in the integration of a self-supervised task that uses another DCRNN as a decoder to predict future EEG signals from the previously cited graph embeddings. This pretraining strategy, combined with specific fine-tuning of the encoding and corresponding classifier block, allows to further improve the outcome of the downstream tasks, when properly fine-tuned.

Multivariate-time-series forecasting is a common task for recurrent GNNs, and recent works proposed methods to refine the graph used by the model in a data-driven way. Shang et al. (2022) present a method, called GST, for network inference to improve the forecasting performances of RGNNs. They learn a parameterized probability distribution over graphs by introducing a differentiable stochastic block. They show some promising experimental results including this block in a DCRNN, the same architecture used by Tang et al. (2021). Xu et al. (2023) extend the idea of GST to produce dynamic networks. Instead of learning a single graph for the full dataset, their parameterized graph-learning block takes as input the corresponding signals, thus leveraging local temporal information.

We propose two architectures for seizure analysis, based respectively on static and dynamic graph learning blocks (Shang et al., 2022; Xu et al., 2023), and study their performances with and without self-supervised pretraining strategies. In Section 2 we describe the model and its building blocks, then in Section 3 we present our experimental pipeline and an in-depth ablation study.

2 METHODS

In this Section we describe the building blocks of our model, a Recurrent Graph Neural Network, whose graph is produced by a learning block which we call Weighted Graph Time Series (*WGTS*).

2.1 DCGRU AND DCRNN

Li et al. (2018) proposed a recurrent GNN, called Diffusion Convolutional Recurrent Neural Network (DCRNN), for forecasting on directed graphs. This architecture captures spatial dependencies using bidirectional random walks on the graph, and temporal dependencies using recurrent units.

DCRNN is a *sequence-to-sequence* model (Sutskever et al., 2014), which predicts the following T steps from a sequence of graph signals of length T':

$$[X^{(t-T'+1)}, \ldots, X^{(t)}; \mathcal{G}] \xmapsto{h(\cdot)} [X^{(t+1)}, \ldots, X^{(t+T)}; \mathcal{G}]. \tag{1}$$

Its building block is *DCGRU*, defined on top of the *Gated Recurrent Unit* (Cho et al., 2014) (GRU) by using *diffusion convolutions* instead of matrix multiplications.

2.2 SELF-SUPERVISED PRETRAINING

Self-supervision is a learning paradigm in which we take a labeled dataset, for which we have a predefined task that we want to solve, and we define one, or more, tasks intrinsic to the structure of the data. The goal is to leverage this structure to refine the embeddings of the model, generally a neural network, in order to get improved performances on the pretext task. More precisely, the model is designed to have an encoding block and multiple "heads", one for each task, that can be trained with their corresponding losses.

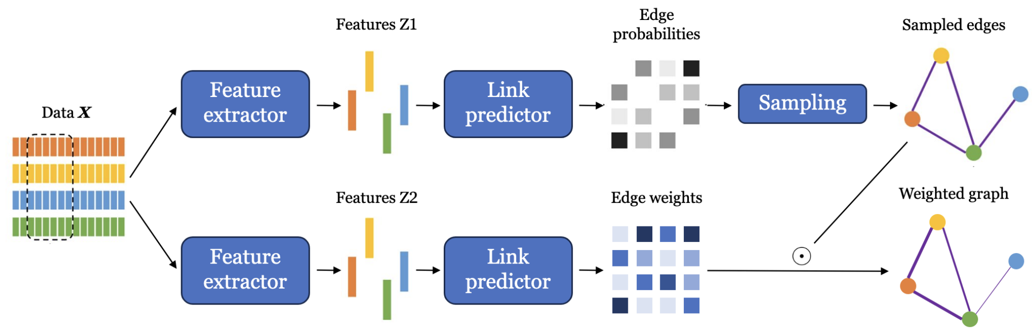

Figure 1: Weighted Graph Time Series (*WGTS*) model. First row illustrates the structure learning and graph sampling schemes, corresponding to basic GTS (Shang et al., 2022). Second row shows the parrallel blocks that produce edge weights, which are assigned to sampled edges. The final weighted graph $\boldsymbol{W}$ is then used in a recurrent GNN such as DCRNN for sequence-to-sequence prediction. Standard *WGTS* feeds all input data to the sequence extractor, while *WGTS Dynamic* receives windowed data, highlighted by the dashed rectangle on the left.

For labeled time series, a natural self-supervised setting is given by forecasting. Classification models generally work with windowed data, so that the encoding block can be shared by a one-step ahead forecasting decoder block. Tang et al. (2021), for instance, propose a three-headed model from which perform either of three downstream tasks, namely graph-signals forecasting (*self-supervised pretraining*), *seizure detection*, or *seizure classification*.

2.3 Graph learning blocks

Although Tang et al. (2021) achieved impressive results with a graph based on the placement of EEG electrodes on the scalp (10-20 EEG graph), we suppose that task-specific networks should provide a better inductive bias to our models.

We build upon the parameterized sampling scheme from Shang et al. (2022) to learn edge probabilities. Their method, called Graph Time Series (GTS), defines a parameterized link predictor, which produces a *structure* matrix $\boldsymbol{P} \in [0,1]^{N\times N}$ from input data. More precisely, a time convolution followed by a fully connected layer produces an embedding $\boldsymbol{z}_i$ for each channel i (*feaure extractor*); then a two-layer perceptron takes the concatenation of two vectors $\boldsymbol{z}_i, \boldsymbol{z}_j$ and gives the corresponding *edge probability* $\boldsymbol{P}_{ij} \in [0,1]$, for all $i, j \in V$ (*link predictor*). They obtain the adjacency matrix $\boldsymbol{W}$ by sampling each edge from a Bernoulli distribution $w_{ij} \sim \mathrm{Ber}(p_{ij})$, and we can back propagate through this sampling by applying the *Gumbel reparameterization trick* (Jang et al., 2017; Maddison et al., 2017).

Xu et al. (2023), extend the GTS block, by applying it to separate windows, instead of the full time series, so that it can produce dynamic networks. We will denote this model as *GTS Dynamic*.

We design a weighted version of these two GTS blocks, which we denote *WGTS*. Figure 1 illustrates our model, which has two parallel structure learning blocks. The first one produces edges probabilities P_t like the original GTS methods, from which we sample the adjacency matrix. The second learning block produces edge weights $\boldsymbol{W}_t$, which we assign to the corresponding edges, if they are sampled.

3 Experiments

We compare our GTS architecture to the following baselines, of which the first two are not based on graphs,

- *LSTM* (Hochreiter & Schmidhuber, 1997), a Long-Short-term memory RNN, with linear layers in its update blocks;

Table 1: Performance comparison of our models (marked by *) and benchmarks on seizure classification (weighted F1 score) and detection (AUROC). For both tasks, higher metrics are better, **best scores** are bold and second best are underlined. We report the average and STD test values obtained across multiple random initializations of each model.

MODEL AND PROPERTIES	CLASSIFICATION F1	DETECTION AUROC
CNN-LSTM	0.617 ±0.023	0.755 ±0.010
LSTM	0.650 ±0.080	0.767 ±0.026
Dist-DCRNN	0.717 ±0.074	0.856 ±0.008
Dist-DCRNN + SSL	0.729 ±0.023	0.860 ±0.015
GTS	0.710 ±0.050	0.856 ±0.020
**WGTS*	0.716 ±0.036	0.822 ±0.031
**WGTS + dist*	0.710 ±0.036	0.845 ±0.016
**WGTS + SSL*	0.762 ±0.013	0.847 ±0.003
**WGTS + SSL + dist*	0.753 ±0.035	0.861 ±0.013
**WGTS Dynamic*	0.737 ±0.062	0.843 ±0.008
**WGTS Dynamic + dist*	0.735 ±0.025	**0.867 ±0.007**
**WGTS Dynamic + SSL*	0.708 ±0.055	0.828 ±0.032
**WGTS Dynamic + SSL + dist*	**0.769 ±0.036**	0.851 ±0.023

- *CNN-LSTM* (Ahmedt-Aristizabal et al., 2020), which shares the LSTM recurrent mechanism, but uses convolutional layers in its update blocks;
- *Dist-DCRNN* (Tang et al., 2021), a recurrent GNN based on the 10-20 EEG distance graph.

We study the influence of multiple factors in our architecture by ablation. In particular, we focus on the following properties

- *SSL*, self-supervised pretraining on one-step ahead forecasting;
- *dist*, we always include the edges of the 10-20 EEG graph on top of the learned one.

Table 1 shows the average performances of all models on both the seizure classification and the seizure detection task, averaged over multiple initializations. We train every architecture on the same standard training set for the TUH dataset, consisting in 592 patients, and compute test scores on the 45 hold-out patients of the test split.

For both classification and detection, the Dynamic Weighted Graph Structure Learning model obtains the best results, significantly outperforming the baselines on unseen patients. We suppose that dynamic networks are more relevant than static ones for seizure analysis, since they might better capture the evolution of anomalous brain activity.

Overall, learning a task-specific graph improves the results when paired with self-supervised pretraining, or with a meaningful prior, such as the *dist* graph. WGTS alone performs on par or worse than baselines, probably due to having many more parameters, but paired with SSL it leverages its overparameterization by boosting test scores by 0.05 for seizure classification, and 0.02 for detection.

4 CONCLUSION

We propose a novel self-supervised graph-learning method that significantly improves the analysis of seizures from brain activity. Our exeriments show that data-driven networks help models to gain a better inductive bias for the task at hand. Nonetheless, the inclusion of relevant priors plays an important role in harvesting the model capabilities.

In our experiments, we mainly focused on structural properties of the data, such as including the geometry of the problem and learning time-varying graphs to capture changing connectivity in epileptic seizures. Still, we envision that by collaborating with neurologists and physicians we could design more relevant priors.

REFERENCES

David Ahmedt-Aristizabal, Tharindu Fernando, Simon Denman, Lars Petersson, Matthew J. Aburn, and Clinton Fookes. Neural Memory Networks for Seizure Type Classification. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 569–575, 2020.

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*, pp. 103–111. Association for Computational Linguistics, 2014.

Robert S. Fisher, Walter van Emde Boas, Warren Blume, Christian Elger, Pierre Genton, Phillip Lee, and Jerome Engel Jr. Epileptic Seizures and Epilepsy: Definitions Proposed by the International League Against Epilepsy (ILAE) and the International Bureau for Epilepsy (IBE). *Epilepsia*, 46 (4):470–472, 2005.

Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8): 1735–1780, 1997.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical Reparameterization with Gumbel-Softmax, 2017.

Kwanhyung Lee, Hyewon Jeong, Seyun Kim, Donghwa Yang, Hoon-Chul Kang, and Edward Choi. Real-Time Seizure Detection using EEG: A Comprehensive Comparison of Recent Approaches under a Realistic Setting, 2022.

Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *International Conference on Learning Representations*, 2018.

Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables, 2017.

Duong Nhu, Mubeen Janmohamed, Ana Antonic-Baker, Piero Perucca, Terence J. O'Brien, Amanda K. Gilligan, Patrick Kwan, Chang Wei Tan, and Levin Kuhlmann. Deep learning for automated epileptiform discharge detection from scalp EEG: A systematic review. *Journal of Neural Engineering*, 19(5):051002, 2022.

Iyad Obeid and Joseph Picone. The Temple University Hospital EEG Data Corpus. *Frontiers in Neuroscience*, 10:196, 2016.

Yannick Roy, Hubert Banville, Isabela Albuquerque, Alexandre Gramfort, Tiago H. Falk, and Jocelyn Faubert. Deep learning-based electroencephalography analysis: A systematic review. *Journal of Neural Engineering*, 16(5):051001, 2019.

Donald L. Schomer and Fernando Lopes da Silva. *Niedermeyer's Electroencephalography: Basic Principles, Clinical Applications, and Related Fields*. Lippincott Williams & Wilkins, 2012. ISBN 978-1-4511-5315-6.

Chao Shang, Jie Chen, and Jinbo Bi. Discrete Graph Structure Learning for Forecasting Multiple Time Series. In *International Conference on Learning Representations*, 2022.

Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to Sequence Learning with Neural Networks, 2014.

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*, 2021.

World Health Organization. Epilepsy: A public health imperative. Technical Report WHO/MSD/MER/19.2, World Health Organization, 2019.

Nancy Xu, Chrysoula Kosma, and Michalis Vazirgiannis. TimeGNN: Temporal Dynamic Graph Learning for Time Series Forecasting, 2023.