Unsupervised seizure localisation with attention-based graph neural networks

Daniele Grattarola^{1,*}, Lorenzo Livi^{2,3}, Cesare Alippi^{1,4}, Richard Wennberg⁵, Taufik Valiante^{6,7,8,9,10}

*For correspondence:
daniele.grattarola@usi.ch (DG)

¹ Faculty of Informatics, Università della Svizzera italiana, Lugano, Switzerland;
² Departments of Computer Science and Mathematics, University of Manitoba, Winnipeg, Canada; ³ Department of Computer Science, University of Exeter, Exeter, United Kingdom; ⁴ Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Milan, Italy; ⁵ Division of Neurology, Department of Medicine, Krembil Brain Institute, Toronto Western Hospital, University of Toronto, Toronto, Canada; ⁶ Department of Surgery, Division of Neurosurgery, University of Toronto, Toronto, Canada; ⁷ Krembil Brain Institute, Division of Clinical and Computational Neuroscience, Toronto, Canada; ⁸ Institute of Medical Sciences, University of Toronto, Toronto, Canada; ⁹ Institute of Biomedical Engineering, University of Toronto, Toronto, Canada; ¹⁰ Electrical and Computer Engineering, University of Toronto, Toronto, Canada

Abstract Graph neural networks (GNNs) and the attention mechanism are two of the most significant advances in artificial intelligence methods over the past few years. The former are neural networks able to process graph-structured data, while the latter learns to selectively focus on those parts of the input that are more relevant for the task at hand. In this paper, we propose a methodology for seizure localisation which combines the two approaches. Our method is composed of several blocks. First, we represent brain states in a compact way by computing functional networks from intracranial electroencephalography recordings, using metrics to quantify the coupling between the activity of different brain areas. Then, we train a GNN to correctly distinguish between functional networks associated with interictal and ictal phases. The GNN is equipped with an attention-based layer which automatically learns to identify those regions of the brain (associated with individual electrodes) that are most important for a correct classification. The localisation of these regions is fully unsupervised, meaning that it does not use any prior information regarding the seizure onset zone. We report results both for human patients and for simulators of brain activity. We show that the regions of interest identified by the GNN strongly correlate with the localisation of the seizure onset zone reported by electroencephalographers. We also show that our GNN exhibits uncertainty on those patients for which the clinical localisation was also unsuccessful, highlighting the robustness of the proposed approach.

Introduction

Epilepsy is a neurological disorder characterised by recurrent episodes of excessive neuronal firing (Stafstrom and Carmant, 2015). In approximately a third of the patients, epilepsy cannot be treated with anti-seizure drugs and resective surgery can be considered as a possible treatment (Kwan and

41 **Brodie, 2000**). The outcome of surgery is crucially dependent on the successful localisation of the
42 seizure onset zone (SOZ) (**Burns et al., 2014; Van Mierlo et al., 2014**).

43 Electroencephalography (EEG) is the mainstay for studying and diagnosing epilepsy, and it is
44 widely used to detect, classify, and localise seizures by recording and processing the electrical ac-
45 tivity of groups of neurons (**Nunez et al., 2006**). However, due to their low spatial resolution, scalp
46 EEG recordings in some cases are not informative enough to successfully localise seizures (**Shah
47 and Mittal, 2014**). In these cases, intracranial EEG recordings (iEEG), in which electrodes are placed
48 directly on or within the brain, provide better spatio-temporal resolution to capture the dynamics
49 of seizure generation and propagation (**Hashiguchi et al., 2007**). However, the high temporal res-
50 olution of iEEG and the complex functional interaction of distant brain areas, especially during
51 seizures, make the interpretation and processing of raw iEEG data a non-trivial task for clinicians.
52 For this reason, a significant branch of epilepsy research is concerned with summarising iEEG data
53 by considering the pairwise (statistical) dependencies between the activity of different brain ar-
54 eas over time (**Van Mierlo et al., 2014**). These dependencies are usually represented by *functional
55 networks* (FNs), in which each node represents a sensor and edges are weighted by a *functional
56 connectivity* (FC) metric (**Bastos and Schoffelen, 2016**).

57 FNs are a widespread tool to study seizure localisation, with early approaches dating back to the
58 1970s (**Gersch and Goddard, 1970; Brazier, 1972**). Seizures have been observed to affect the func-
59 tional organisation of brain activity at the meso-scale, both from a node-centric (**Burns et al., 2014**)
60 and an edge-centric (**Khambhati et al., 2015**) perspective. In particular, **Burns et al. (2014)** identified
61 sets of brain states that emerge by clustering FNs, consistent in interictal and ictal periods for in-
62 dividual patients. They observed that changes in node centrality in FNs accurately predict the SOZ.
63 **Khambhati et al. (2015)** observed a strengthening of FC in the SOZ during seizures, also coinciding
64 with a topological tightening of the connections (*i.e.*, strong connections also become physically
65 closer). **Khambhati et al. (2016)** proposed *virtual cortical resection*, *i.e.*, the removal of nodes from
66 FNs, in order to study changes in network synchronizability, which is a known predictor for the
67 spread of seizures (**Schindler et al., 2008**). **Lopes et al. (2017)** also observed that the resection of
68 brain areas associated to *rich-club* hubs in FNs is correlated with a good post-operative outcome.
69 Seizure localisation has also been studied in FNs obtained from functional magnetic resonance
70 imaging (fMRI) (**Lee et al., 2014; Weaver et al., 2013**) and scalp EEG (**Staljanse et al., 2017**) data.
71 Recent work by **Covert et al. (2019)** introduced the use of spatio-temporal graph convolutional net-
72 works (ST-GCNs) (**Yu et al., 2017**) to perform seizure detection, and conducted an *ex post* analysis
73 similar to the one of **Khambhati et al. (2016)** to quantify the importance of a node by observing
74 the effect of its removal on the downstream detection accuracy. **Gadgil et al. (2020)** also proposed
75 a methodology based on ST-GCNs that allows identifying high-interest areas in fMRI by learning
76 to estimate edge importance, although they did not apply it to seizure localisation. For a more
77 in-depth review of approaches to seizure localisation with FNs, we refer the reader to **Van Mierlo
78 et al. (2014)**.

79 The aim of this paper is to use the representation of brain states as FNs in order to automate
80 the localisation of seizures using deep learning. Advances in deep learning techniques over the
81 past decade have revolutionised how high-dimensional, high-volume data can be used in the con-
82 text of artificially intelligent systems. In particular, deep learning techniques for computer vision
83 have shown how artificial intelligence can be successfully adopted in clinical settings to aid human
84 experts in their decision making (**Litjens et al., 2017**). Despite these successes, traditional deep
85 learning methods are limited to processing regular structures like images and time series, and
86 cannot naturally consider the relations that exist in a complex system with multiple interacting
87 components, such as those described by FNs evolving over time. For this reason, recent literature
88 has seen the rise of Graph Neural Networks (GNNs) (**Battaglia et al., 2018; Bronstein et al., 2017**)
89 as a generalisation of deep learning techniques to process data represented as arbitrary graphs.

90 In this paper, we introduce a GNN-based methodology for seizure localisation, using FNs to effi-
91 ciently represent brain states. The core of our algorithm is a GNN equipped with an *attention-based*

Figure 1. Schematic view of our GNN-based pipeline for seizure detection and localisation. Starting from raw iEEG data, we compute a functional network to represent the spatio-temporal dynamics of the signals in a compact way. The FN is then given as input to a GNN composed of an edge-aware message passing operation followed by an attention-based readout to compute a graph-level embedding. The embedding is then classified to perform seizure detection, while the attention scores are analysed to perform seizure localisation.

readout. By training such a GNN to perform seizure detection, the readout automatically learns to assign a higher attention to those nodes that are more important for a correct classification. Then, we propose a simple and fast way of analysing the attention coefficients over time, so that we obtain a ranking of the nodes based on their overall importance in detecting a seizure. Crucially, our methodology does not require *a priori* information regarding the SOZ, but only weak supervision in the form of annotated seizure onsets and offsets. The localisation procedure is, therefore, unsupervised. A schematic representation of our approach is shown in Figure 1.

We validate the proposed methodology on clinical iEEG data collected from eight human subjects and show that the electrode rankings computed with our localisation procedure are highly correlated with the true SOZs. We also validate our algorithm on simulated data, using a simple model of seizure initiation (Benjamin et al., 2012) and a more complex brain simulator (Sanz Leon et al., 2013) based on the Epileptor model (Jirsa et al., 2014). Our main contributions and results are summarised as follows:

- We present a new algorithm for unsupervised seizure localisation based on GNNs, which uses FNs to represent brain states in a compact form and requires no explicit supervision on the SOZ;
- We show that the attention coefficients learned by the GNN correlate with clinically-identified SOZs and accurately predict the presence of ictal activity;
- We show that, when electroencephalographers were not able to identify the SOZ from the iEEG data, the GNN also shows uncertainty in the localisation;
- We show that, as expected, the choice of FC metric used to estimate FNs is important for an accurate localisation;
- Finally, we show that our methodology performs well on very imbalanced datasets, achieving a good localisation accuracy even on patients for which we observe as few as five seizures during training.

Methods

Notation. We denote a time series $x_i(t)$ to represent the i -th iEEG channel at time t . We define a graph as a tuple $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_N\}$ represents the set of attributed nodes with attributes $\mathbf{v}_i \in \mathbb{R}^F$, and $\mathcal{E} = \{\mathbf{e}_{i \rightarrow j} | \mathbf{v}_i, \mathbf{v}_j \in \mathcal{V}\}$ represents the set of attributed edges with attributes $\mathbf{e}_{i \rightarrow j} \in \mathbb{R}^S$ indicating a directed edge between the i -th and the j -th node. We indicate the neighbourhood of node i with $\mathcal{N}(i) = \{\mathbf{v}_k | \mathbf{e}_{k \rightarrow i} \in \mathcal{E}\}$. We say that a graph is undirected if $\mathbf{e}_{i \rightarrow j} \in \mathcal{E} \iff \mathbf{e}_{j \rightarrow i} \in \mathcal{E}$. Note that in the text, for simplicity, we refer to nodes by means of their index, e.g., node i .

Functional networks

Choosing a suitable FC metric to model the pairwise interaction between brain areas is a non-trivial challenge, as there exist a large variety of methods with their own advantages and disadvantages. FC metrics can be characterised according to several properties, including whether they are in the time or frequency domain, whether they are directed or undirected (*i.e.*, if they model asymmetric or symmetric couplings), or whether they are model-free or model-based (**Bastos and Schoffelen, 2016**). Here, we focus on undirected FC metrics to simplify the GNN computation, and on model-based approaches to reduce the computational costs of estimating the FC metrics directly from data. We do, however, consider two different metrics to highlight the practical differences that emerge between time- and frequency-domain metrics.

FNs are generated by computing a FC value for each pair of iEEG channels $x_a(t)$ and $x_b(t)$ over a time window of length T . For the time-domain metric, we consider Pearson's correlation coefficient:

$$\mathbf{e}_{a \rightarrow b} = \mathbf{e}_{b \rightarrow a} = \frac{\sum_{t=1}^T (x_a(t) - \bar{x}_a)(x_b(t) - \bar{x}_b)}{\sqrt{\sum_{t=1}^T (x_a(t) - \bar{x}_a)^2} \sqrt{\sum_{t=1}^T (x_b(t) - \bar{x}_b)^2}}, \quad (1)$$

where $\bar{x}_a = \frac{1}{T} \sum_{t=1}^T x_a(t)$ and analogously for $\bar{x}_b$. Correlation allows to quantify symmetric linear interactions, it is easy to compute and, as such, it is often used in the literature. For the frequency domain, we consider the phase-locking value (PLV) (**Lachaux et al., 1999**):

$$\mathbf{e}_{a \rightarrow b} = \mathbf{e}_{b \rightarrow a} = \left| \frac{1}{T} \sum_{t=1}^T e^{i(\varphi_a(t) - \varphi_b(t))} \right|, \quad (2)$$

where $\varphi_a(t)$ indicates the instantaneous phase of signal $x_a(t)$ obtained via Hilbert transform (and similarly for $\varphi_b(t)$). A significant advantage of PLV over correlation is that it is less sensitive to artifacts in the iEEG signals (such as those caused by the patient's movements). After computing the FC metrics for each pair of channels, we sparsify the resulting FNs by removing those edges for which $|\mathbf{e}_{i \rightarrow j}| < 0.1$, *i.e.*, those indicating weak coupling.

We generate a dataset of FNs for each patient, dividing the FNs into ictal and interictal classes and proceeding in a per-seizure fashion. Let f_s be the sampling rate of the iEEG signal, L the duration of a seizure, t_0 the time indicating the seizure onset, $k \geq 1$ a subsampling factor, and T the length of the time windows. Additionally, let $y(t) \in \{0, 1\}$ be a binary signal indicating whether the patient is having a seizure at time t (*i.e.*, $y(t) = 1$ if $t \geq t_0$ and 0 otherwise). Note that we consider each seizure to end at time $t_0 + L$ and we do not compute FNs for the data immediately following a seizure offset.

Given a time window $[t - T, \dots, t]$, we compute a FN $\mathcal{G}^{(t)}$ and label it with class

$$\mathcal{Y}^{(t)} = \begin{cases} 1, & \text{if } \sum_{\tau=t-T}^t y(\tau) > T/2 \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

To generate the FNs associated with seizures (class 1), we consider the data interval $[t_0 - T/2, \dots, t_0 + L]$ and take overlapping windows of size T with a stride of $1/f_s$. For the interictal FNs (class 0), instead, we consider a longer period preceding the seizure onset, $[t_0 - kL, \dots, t_0 + T/2]$, and we take windows at a larger stride of k/f_s . In this work, we consider $k = 10$ and $T = 1$ s for all experiments, although other values are possible.

This procedure to generate the FNs (summarised in Figure 2) results in a balanced dataset and has two advantages. First, it allows us to fully use all the available (and rare) ictal events. Second, it allows us to consider a more diverse sample for the interictal class. The small differences between consecutive FNs of the positive class, due to the small stride at which windows are taken, can be seen as a form of sample weighting to account for the class unbalance characterising the problem.

Figure 2. Schematic representation of the procedure used to generate FNs. For each seizure of length L starting at t_0 (marked in red), we consider an interictal interval of length kL . Interictal FNs are generated taking windows of length T at stride k/f_s , while ictal windows are taken with stride $1/f_s$ (in green). For each window and each pair of electrodes i and j , we compute the FC value $e_{i \rightarrow j}$ (in blue) to obtain the full FN. This figure is only meant to represent the procedure and is not shown in any physical temporal scale.

In order to have initial node features that can be processed by the GNN, we consider dummy attributes set to 1 for all nodes. Other choices that depend on the actual iEEG signals are possible (e.g., the signal power or wavelet coefficients) but were not explored in this work.

Attention mechanism

Attention (Bahdanau et al., 2014; Vaswani et al., 2017) is a processing technique for neural networks to learn how to selectively focus on parts of the input. Originally developed for aligning sentences in neural machine translation (Bahdanau et al., 2014; Vaswani et al., 2017), the attention mechanism has been used to achieve state-of-the-art results on different tasks like language modelling (Brown et al., 2020), image processing (Xu et al., 2015), and even learning on graphs (Velickovic et al., 2018).

In this paper, we focus on the concept of *self*-attention, which indicates a class of attention mechanisms that learn to attend to the output of a layer using the output itself (in contrast to classical attention, which uses the output of one layer to focus on the output of another – e.g., the sentence of the source language is used to focus on the target language). At its core, self-attention consists of computing a compatibility score $\alpha_{ij} \in [0, 1]$ between two vectors $\mathbf{h}_i, \mathbf{h}_j \in \mathbb{R}^F$ (both part of the same sequence, image, graph, etc.):

$$\alpha_{ij} = \text{Softmax}_j(e_{ij}) = \frac{\exp(e_{ij})}{\sum_{k=1}^N \exp(e_{ik})}, \quad (4)$$

where

$$e_{ij} = a(\mathbf{h}_i, \mathbf{h}_j) \quad (5)$$

and a is called an *alignment* model, which is usually learned end-to-end along with the other parameters of the neural network. The compatibility score is then used to compute a representation of element i as:

$$\mathbf{z}_i = \sum_j \alpha_{ij} \mathbf{h}_j. \quad (6)$$

Intuitively, the attention mechanism learns the importance of element j to describe element i , and computes score α_{ij} to quantify this importance. The alignment model can be seen as a similarity function between the two elements, which is then normalised via the Softmax function. Different implementations of the alignment model are possible, although often it is implemented as a multi-layer perceptron.

Attention mechanisms are usually trained without direct supervision and automatically learn to focus on different parts of the data according to the loss of the given task. By optimising the overall task loss, the attention layers in a neural network learn to compute the optimal compatibility scores.

191 This is a key aspect of our proposed methodology, where we use self-attention to automatically
192 detect those brain areas (monitored via different iEEG channels) that are important to detect a
193 seizure.

194 **Graph neural networks for seizure localisation**

195 Graph Neural Networks (GNNs) are a class of neural networks designed to perform inference on
196 graph-structured data (Battaglia et al., 2018). At their core, GNNs learn to represent the nodes
197 of a graph by propagating information between connected neighbours, whereas a global repre-
198 sentation of the entire graph is usually obtained by computing a *readout* of the nodes, like a sum,
199 average, or component-wise maximum vector. In this work, we focus on the family of *message-*
200 *passing* networks (Gilmer et al., 2017), in which the l -th layer maps the attributes $\mathbf{h}_i^{(l-1)} \in \mathbb{R}^{F^{(l-1)}}$ of
201 the i -th node to:

$$\mathbf{h}_i^{(l)} = \gamma \left(\mathbf{h}_i^{(l-1)}, \square_{j \in \mathcal{N}(i)} \phi \left(\mathbf{h}_i^{(l-1)}, \mathbf{h}_j^{(l-1)}, \mathbf{e}_{j \rightarrow i} \right) \right), \quad (7)$$

202 where $\mathbf{h}_i^{(l)} \in \mathbb{R}^{F^{(l)}}$, $\mathbf{h}_i^{(0)} = \mathbf{v}_i$, and ϕ and γ are differentiable functions equivariant to node permu-
203 tations, respectively called the *message* and *update* functions, while $\square$ is a permutation-invariant
204 function (such as the sum or the average) to aggregate incoming messages.

205 Many recent papers have introduced methods for graph representation learning based on this
206 general scheme, with different implementations ranging from polynomial (Defferrard et al., 2016)
207 or rational (Bianchi et al., 2019) graph convolutional filters, to attentional mechanisms (Velickovic
208 et al., 2018). In most of these works the creation of messages is only dependent on the node
209 attributes, although some methods have been proposed that also explicitly take edge attributes
210 into account (Simonovsky and Komodakis, 2017; Schlichtkrull et al., 2018). In particular, the Edge-
211 Conditioned Convolutional (ECC) operator proposed by Simonovsky and Komodakis (Simonovsky
212 and Komodakis, 2017) incorporates edge attributes into the message-passing scheme by using a
213 *kernel-generating network* $f^{(l)}(\cdot)$ that dynamically computes messages between each pair of con-
214 nected nodes. An ECC layer is thus defined as:

$$\mathbf{h}_i^{(l)} = \mathbf{h}_i^{(l-1)} \cdot \mathbf{W}_{\text{root}}^{(l)} + \sum_{j \in \mathcal{N}(i)} \mathbf{h}_j^{(l-1)} \cdot f^{(l)}(\mathbf{e}_{j \rightarrow i}), \quad (8)$$

215 where $\mathbf{W}_{\text{root}}^{(l)} \in \mathbb{R}^{F^{(l-1)} \times F^{(l)}}$ is a learnable kernel applied to the root node itself and the kernel-generating
216 network is usually a multi-layer perceptron $f^{(l)} : \mathbb{R}^S \rightarrow \mathbb{R}^{F^{(l-1)} \times F^{(l)}}$.

217 Our method for seizure localisation can be summarised as follows. First, we train a GNN with
218 an attention-based readout to detect seizures from FNs. This is a graph-level classification problem
219 where a label (ictal or interictal) is assigned to each FN. Then, we analyse the compatibility scores
220 learned by the attentional mechanism to identify those nodes that the model consistently consid-
221 ers as important. Although we train the GNN to do seizure detection in a supervised way, *i.e.*, it
222 requires manually-annotated seizure onsets and offsets, the localisation is fully unsupervised. This
223 is one of the main strengths of the proposed method, as significantly less manual work is required
224 to annotate the temporal boundary for each individual seizure, rather than the SOZ.

225 There are two main components in our GNN architecture. First, the connectivity information
226 is propagated to the node attributes via an edge-aware message-passing operation like ECC. A
227 single layer is sufficient because the input FNs are densely connected, and most nodes will receive
228 information from the whole graph in a single step of message passing.

229 Then, we use a self-attentional mechanism to compute the graph readout:

$$\mathbf{z} = \text{Attn-RO}(\mathbf{h}) = \sum_{j=1}^N \alpha_j \mathbf{h}_j \quad (9)$$

230 where

$$\alpha_j = \frac{\exp(\mathbf{h}_j \cdot \mathbf{a})}{\sum_{k=1}^N \exp(\mathbf{h}_k \cdot \mathbf{a})}, \quad (10)$$

231 $\mathbf{h}_j \in \mathbb{R}^{F^{out}}$ is the embedding of the j -th node computed by the ECC layer, and $\mathbf{a} \in \mathbb{R}^{F^{out}}$ is a vector of
 232 learnable weights. Note that, compared to Equation (6), here index i is left implicit as the attention
 233 is only computed once for all nodes, to reduce the graph to a vector. This is also reflected in the
 234 fact that the alignment model is a function of only one node at a time, e.g., $\mathbf{h}_j \cdot \mathbf{a}$. For a more general
 235 way of applying attention to every possible pair of nodes (while maintaining the graph structure),
 236 see (Velickovic et al., 2018).

237 Finally, a multi-layer perceptron $\text{MLP}(\cdot)$ with sigmoid activation computes the probability that
 238 the input FN represents an ictal window of iEEG data.

239 The full architecture is written as:

$$\hat{y} = \text{MLP}(\text{Attn-RO}(\text{ECC}(\mathcal{G}))) \quad (11)$$

240 where $\mathcal{G}$ represents an input FN (cf. Figure 1).

241 By training the GNN to correctly distinguish the ictal FNs from the non-ictal ones, we also implic-
 242 itly train the attentional readout Attn-RO to assign higher attention to those nodes of the FNs that
 243 maximise the confidence in the prediction. We then analyse how the attention scores assigned
 244 to nodes change over time, and rank the nodes according to the overall amount of attention that
 245 they receive before and during a seizure. The localisation procedure is described in the following
 246 section.

247 Localising the seizure onset zone

248 For each seizure in the data, we consider symmetric intervals of length $2L$ centred at the seizure
 249 onset, so that the first L timesteps are pre-ictal and the remaining L cover the beginning of the
 250 seizure. For each of the $2L$ timesteps, we compute a FN $\mathcal{G}^{(t)}$ from a $T = 1$ s window ending at time
 251 t , obtaining a sequence of FNs $[\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(2L)}]$ (this is equivalent to how we generate the training
 252 datasets, except that the subsampling is set at $k = 1$). For each FN in the sequence, we use the
 253 GNN to compute the attention scores over the nodes according to Equation (10). We thus compute
 254 a sequence of attention scores $[\alpha_i^{(1)}, \dots, \alpha_i^{(2L)}]$ for each node i .

255 We then sum the sequence of attention scores to obtain the overall *importance* of the node
 256 over the considered time interval:

$$\sigma_i = \sum_{t=1}^{2L} \alpha_i^{(t)}, \quad (12)$$

257 and normalise the importance scores to the $[0, 1]$ interval as:

$$s_i^{(s)} = \frac{\sigma_i^{(s)} - \min_{j \in \mathcal{V}} \sigma_j^{(s)}}{\max_{j \in \mathcal{V}} \sigma_j^{(s)} - \min_{j \in \mathcal{V}} \sigma_j^{(s)}}. \quad (13)$$

258 Finally, we rank the nodes according to their importance and predict the SOZ accordingly.

259 Results

260 We report the results obtained on real iEEG data collected from eight patients. Additional results on
 261 two brain activity simulators (a simple network model (Benjamin et al., 2012) and *The Virtual Brain*
 262 simulator (Sanz Leon et al., 2013)) and all experimental details regarding the GNN are reported in
 263 the appendix.

264 Data collection and pre-processing

265 We used iEEG data recorded from eight human subjects with medically refractory epilepsy, the
 266 recordings obtained as part of their standard clinical pre-surgical investigations. The study was
 267 approved by the Research Ethics Board at the University Health Network (ID number 12-0413) and
 268 written consent for data collection was obtained from all participants. Each patient had a varying
 269 number of recorded clinical seizures and the number of electrodes also varied from patient to
 270 patient (cf. Table 1). The data was recorded from subdural or intracerebral depth electrodes at

Table 1. Summary of the patients considered for this study. The columns indicate (left-to-right): the number of recorded seizures, the number of implanted electrodes, the presence of ictal activity (IA) marked by electroencephalographers on one or more channels, whether the patient had surgery, and the outcome of the surgery.

Patient	Seizures	Electrodes	IA identified	Surgery	Outcome
66	15	100	Yes, low confidence	No	-
70	9	96	Yes	Yes	Seizures reduced
75	10	23	Yes	No	-
76	5	74	No	No	-
77	11	38	Yes	Yes	Seizures reduced
78	18	45	Yes, poorly defined	No	-
85	5	45	Yes	No	-
87	16	69	Yes	No	-

$f_s = 500\text{Hz}$ over the course of several days per patient, and seizures were manually annotated by electroencephalographers, inspecting both raw iEEG and video recordings of the patient. The iEEG signal was notch-filtered at 60Hz and related harmonics to remove powerline trends, and then filtered with an order-3 low-pass filter at 100Hz to remove any high-frequency noise. Then, each electrode channel was independently re-referenced to have zero mean and rescaled to have unit variance.

Before pre-processing, we visually inspected the raw data of each patient and each seizure to assess the presence of bad channels: we considered symmetric windows around each labelled seizure onset and we removed from the data any channels that exhibited abnormal (*i.e.*, either flat or excessive) activity in at least one seizure.

Per-patient analysis of the SOZ

This section reports the available clinical data for the patients considered in our study. For all patients, both the seizure onset time instants and the SOZ annotations were provided by electroencephalographers.

Patient 66 demonstrated ictal activity in both the left and right posterior interhemispheric regions (Figure 3a), with interictal epileptiform discharges recorded independently from the left anterior frontal and right middle frontal lobes. The patient did not undergo resective surgery due to a low confidence in the identification of the SOZ. Patient 70 showed clear seizures originating in the right posterior insular region (Figure 3b). The patient underwent laser interstitial thermal therapy targeting a focal cortical dysplasia in the area. The patient continued to have some post-operative seizures, although these were reduced in frequency and intensity, indicating that the SOZ was identified correctly. Patient 75 had seizure onsets recorded independently from both temporal lobes and thus was not a candidate for surgery. Patient 76 had no clear ictal activity identified by electroencephalographers in the iEEG recordings and was thus not a candidate for surgery, the SOZ evidently not captured by the intracranial electrode placements. Patient 77 demonstrated ictal activity in the left hippocampal body, and underwent a left anterior temporal resection. The patient continued to have seizures after the surgery, but of reduced frequency and intensity, indicating a successful localisation of the SOZ. Patient 78 had multiple seizures recorded with poorly defined, inconsistent ictal onsets over temporoparietal sensory cortex and was deemed not a candidate for surgical resection due to uncertainty on the SOZ. Patient 85 had seizures recorded in the left hemisphere, with onsets involving a broad region of temporal lobe neocortex. The patient was not subject to resection due to the epileptogenic zone being too large, and near eloquent language cortex. Patient 87 exhibited abnormal activity in the left amygdala and hippocampus. The patient had already undergone contralateral right anterior temporal resective surgery years prior to the

Figure 3. Examples of raw iEEG traces for patients 66 and 70. The two plots show the activity of electrodes that were identified as SOZs by electroencephalographers. The vertical line marks the seizure onset, as reported in the patients' clinical records.

collection of the iEEG data and was not a candidate for further resections.

Table 1 summarises the relevant details of the eight patients. In particular, six patients had clinically-identified, well-defined information regarding the SOZ, whereas in two patients the SOZ could not be clearly identified in the iEEG data by electroencephalographers.

Results on seizure detection and localisation

Table 2 reports the average Area Under the Receiver Operating Characteristic Curve (ROC-AUC) and the Area Under the Precision-Recall Curve (PR-AUC) obtained by the GNN on the seizure detection task. We report the results obtained using both FC metrics (correlation and PLV) to generate the FNs. We also report the detection performance of a baseline convolutional neural network for time series classification (details in the Appendix).

The GNN achieved an average ROC-AUC score of 79.56 and an average PR-AUC of 81.24 (the average is computed over all patients) when using correlation as FC metric. These results are aligned with the performance of the baseline, which our method slightly outperformed on average, and indicate that 1) our choice of architecture was reasonable and 2) using graph-structured data is an interesting direction for future research on efficient seizure detection. We also recall that the detection task is only meant to provide a weak supervision for the more interesting challenge of localisation, and that better detection results could be achieved by increasing the capacity of the GNN or collecting more training data.

Tables 3 and 4 report the performance of the model on the patients with a known SOZ, respectively using correlation and PLV to generate FNs. In particular, we report three main performance measures:

- (a) the average precision at K (AP@ K) (Sanderson et al., 2010) obtained by the GNN when computing an average ranking of the electrodes. Each electrode is re-ranked by considering five models trained on the same data and taking the average score assigned to each electrode over all models and all seizures. This measure quantifies the GNN's ability to correctly identify

Table 2. Average ROC-AUC score and average PR-AUC score for seizure detection on unseen test data. These scores represent the model's ability to correctly classify the FNs as interictal or ictal. The last row reports the average score over all patients. The highest ROC-AUC and PR-AUC scores are reported in bold for each patient.

Patient	Baseline		GNN Corr.		GNN PLV	
	ROC	PR	ROC	PR	ROC	PR
66	62.54 $\pm$ 22.5	70.06 $\pm$ 17.8	68.63 $\pm$ 11.43	75.20 $\pm$ 10.30	75.68 $\pm$ 23.3	77.51 $\pm$ 20.1
70	80.19 $\pm$ 15.5	85.96 $\pm$ 10.6	86.87 $\pm$ 9.07	89.04 $\pm$ 9.35	65.36 $\pm$ 20.1	72.91 $\pm$ 14.8
75	82.32 $\pm$ 14.19	87.25 $\pm$ 9.24	93.35 $\pm$ 3.12	94.34 $\pm$ 2.72	71.50 $\pm$ 14.8	71.02 $\pm$ 16.3
76	67.81 $\pm$ 8.75	69.83 $\pm$ 13.12	60.40 $\pm$ 14.41	61.11 $\pm$ 14.82	53.83 $\pm$ 6.6	51.67 $\pm$ 6.4
77	76.18 $\pm$ 15.41	80.42 $\pm$ 14.26	77.04 $\pm$ 11.98	76.39 $\pm$ 13.03	71.46 $\pm$ 12.1	71.45 $\pm$ 12.9
78	76.32 $\pm$ 17.2	80.94 $\pm$ 13.5	73.72 $\pm$ 17.14	76.02 $\pm$ 14.53	63.81 $\pm$ 17.2	71.06 $\pm$ 12.4
85	76.46 $\pm$ 11.24	81.22 $\pm$ 7.65	85.52 $\pm$ 10.95	85.92 $\pm$ 13.65	69.32 $\pm$ 2.6	65.55 $\pm$ 1.8
87	85.60 $\pm$ 14.6	89.29 $\pm$ 10.7	90.97 $\pm$ 5.51	91.89 $\pm$ 3.49	77.69 $\pm$ 11.5	78.32 $\pm$ 11.3
Avg.	75.93 $\pm$ 7.06	80.62 $\pm$ 6.86	79.56 $\pm$ 10.82	81.24 $\pm$ 10.37	68.58 $\pm$ 7.08	69.94 $\pm$ 7.86

- the SOZ for a patient in general, which is the most clinically relevant scenario.
- (b) The mean AP@ K (MAP@ K) obtained by the GNN on different individual seizures. In this case, the ranking for each seizure is compared to the ground truth independently of the others (*i.e.*, without averaging the scores), and the scores are averaged *a posteriori* (also considering five repetitions of the experiments). This measure quantifies the GNN's ability to correctly identify target electrodes in a given seizure.
- (c) The MAP@ K obtained by the GNN on different individual seizures, but considering groups of electrodes belonging to the same strip (implying spatial locality of the electrodes). This allows us to evaluate the performance of the model at a coarser scale.

From the results we see that, while correlation was a clearly better metric for the task of seizure detection, the localisation performance can vary depending on the particular FC metric used. In particular, the localisation for patients 66 and 77 was better when using correlation networks, but PLV yielded better results for patients 75, 85, and 87.

In general, however, we note that the (M)AP@5 score is positive for both FC metrics, for all performance measures and for all patients, meaning that at least one SOZ-associated electrode was ranked in the top five every time. We also note that the GNN achieves a perfect AP@2 score (average rankings) in six out of eight cases when using PLV, indicating a high chance of localising at least two relevant electrodes per patient.

Remarkably, we see that these results were obtained even when considering small datasets, *e.g.*, down to only five seizures for patient 85 (*cf.* Table 1).

Comparison with clinical information

Figure 5 summarises our results and provides an overview of the importance scores, their variability across different models and seizures, and their agreement with the ground truth.

The results for patient 77 can be considered a complete success, with the highest AP@ K scores among all patients and very little uncertainty in the ranking by the GNN. Crucially, the successful post-operative outcome confirms that the localisation of the SOZ for this patient was accurate and points to a strong localisation ability of the GNN. For patient 70, ictal activity was evident and well-localised on a specific depth electrode placed in the right insular complex (RINS1). The clinical localisation of the SOZ was therefore likely accurate, even if the outcome of the surgery was not completely successful. More importantly, we notice that the GNN was strongly aligned with the human analysis given the same information, and similarly focused on the same electrode (which

Table 3. Localisation performance for patients with a known SOZ, when using Pearson's correlation as FC metric. We report: **(a)** the average precision at K for averaged rankings, which evaluates the localisation for the patient overall; **(b)** the mean average precision at K for single rankings, which evaluates the localisation for a given seizure; **(c)** the mean average precision at K for single rankings and groups of electrodes, which is equivalent to (b) but at a coarser scale. We report scores for $K = 2, 5, 10$. Bold indicates that the results are better than the ones obtained with PLV as FC metric (*cf.* Table 4).

Patient	(a) AP@ K - Avg. rank			(b) MAP@ K - Single			(c) MAP@ K - Groups		
	$K = 2$	$K = 5$	$K = 10$	$K = 2$	$K = 5$	$K = 10$	$K = 2$	$K = 5$	$K = 10$
66	50.00	20.00	12.50	22.31	12.0	7.24	26.92	21.48	31.64
70	100.00	100.00	100.00	51.11	54.8	56.71	53.33	58.48	60.73
75	0.00	16.67	38.96	20.37	26.51	28.98	36.11	45.09	50.07
77	100.00	55.00	55.00	97.73	48.55	54.71	99.09	99.09	99.09
85	0.00	0.00	0.00	0.00	2.71	5.43	25.00	53.75	64.42
87	0.00	6.67	5.56	19.69	13.00	7.42	20.00	36.43	44.07

Table 4. Localisation performance for patients with a known SOZ, when using PLV as FC metric. Bold indicates that the results are better than the ones obtained with correlation as FC metric (*cf.* Table 3).

Patient	(a) AP@ K - Avg. rank			(b) MAP@ K - Single			(c) MAP@ K - Groups		
	$K = 2$	$K = 5$	$K = 10$	$K = 2$	$K = 5$	$K = 10$	$K = 2$	$K = 5$	$K = 10$
66	0.00	5.00	5.83	8.67	5.97	4.67	16.67	18.82	31.78
70	100.00	100.00	100.00	50.00	54.33	56.65	50.00	58.04	60.21
75	100.00	55.00	45.46	60.00	40.82	32.58	66.88	45.16	51.36
77	100.00	40.00	48.57	66.82	38.28	45.27	91.82	93.48	93.48
85	100.00	40.00	28.29	70.00	40.40	25.43	66.00	65.02	78.35
87	50.00	20.00	10.00	15.62	9.15	6.43	16.56	20.07	32.30

Figure 4. Histograms of the attention scores over a 2-second window starting from a seizure onset. Each bin represents the frequency with which the corresponding attention score is assigned to ten randomly-selected electrodes. Figure (a) shows a patient with a known SOZ, while Figures (b) and (c) show patients without a known SOZ. For Figure (a), the contribution to each bin of those electrodes that are part of the SOZ ground truth are highlighted in orange. Note how the score distribution for SOZ-associated electrodes is spread out towards higher values, while for patients with no known SOZ the scores are similar for all channels.

is ranked first using either of the FC metrics). Our methodology also confirms the conclusions reached by electroencephalographers for patients 75, 85 and 87, although further studies would be required to give a more precise interpretation of the results (including, possibly, the outcome of future surgeries). The results for patient 87 are particularly uncertain, despite the GNN achieving a good detection accuracy (*cf.* Table 2). In general, however, the rankings provided by the GNN show a high agreement with the medical assessment in those cases where the SOZ was successfully identified.

For patients with no known SOZ (76, 78) the GNN has a low detection performance and the average attention scores assigned by the GNN are uniformly distributed across all electrodes around an average score of 0.5. On the contrary, patients with a known SOZ have a few electrodes that are assigned a majority of the attentional budget. This difference between the two cases is more clearly visualised in Figure 4, which shows the distribution of the scores given to different electrodes at the seizure onset (patient 77 is taken as representative of the case in which the SOZ is known).

For patient 66, the GNN did not identify any particularly important regions despite there being some clinical evidence of ictal activity in the posterior interhemispheric region. Two posterior interhemispheric electrodes are indeed ranked in the top ten (averaged rankings) by the GNN when using correlation FNs, although with a very high uncertainty. We note, however, that the uncertainty showed by the GNN was also reflected clinically in the electroencephalographers' interpretations and in the final decision to not operate this patient.

Our analysis for patients 66, 76, and 78 shows that the uncertainty of the GNN correlates with uncertainty or inability on the part of electroencephalographers to identify the SOZ in iEEG, and can still be useful to support their decision making (*e.g.*, deciding to not operate a patient can be just as valuable as a successful localisation).

Discussion

Our work introduces a methodology for automated and unsupervised seizure localisation using graph-based machine learning. Our approach does not require any manual annotation of the SOZ in order to work, making it cheaper to train and easier to scale to a larger number of patients. Our method is also data-efficient: we were able to provide a good – and clinically verified – localisation using as little as five annotated seizures per patient.

The goal of the proposed approach is to provide a support tool for clinicians to allocate precious resources in the analysis of iEEG data, and to improve the efficiency of the decision-making

Figure 5. Top ten electrodes when considering the averaged rankings. We report the ranking obtained with the best-performing FC metric for each patient, according to the AP@10 score for average rankings reported in Tables 3 and 4. The two plots in red indicate those patients for which the SOZ was not identified clinically. Bold labels indicate that the corresponding electrode was marked as a potential SOZ by electroencephalographers.

process. Crucially, in this regard, we note that our algorithm is conservative in scoring potential SOZ candidates. When the SOZ was not identifiable by electroencephalographers, the GNN also showed uncertainty in the scoring (rather than making high-confidence predictions). Contrarily, a high importance score consistently correlated with clinically-identified SOZs. With this premise, we believe that our approach could have practical value if deployed to epilepsy monitoring units to provide a real-time analysis of iEEG recordings.

Future work

There are several directions for future research that could stem from this work. First, we note that by 1) increasing the capacity of the network (in terms of parameters and depth), 2) performing a patient-specific hyperparameter search, and 3) having more seizures on which to train the model, it is likely that both the detection and localisation performance would significantly improve. Also, a possible extension of the proposed methodology could be to explicitly introduce a supervised objective to train the attentional readout using the available information on the SOZ. This would require a per-seizure annotation of every electrode (or, even better, an annotation over time), but could lead to a more accurate localisation. An interesting application of this methodology could also be to provide a patient-agnostic localisation, by training the GNN concurrently on seizures of different patients.

Future work could also explore more in-depth the use of different, combined FC metrics and their impact on the detection and localisation performance. In fact, our results show that this choice could lead to significantly different outcomes (as we showed with correlation and PLV). A way to identify *a priori* the best FC metric to build FNs for a specific patient could bring significant benefits.

Conclusion

We presented a methodology for unsupervised seizure localisation based on GNNs with an attention mechanism. Our approach takes advantage of a compact representation of brain states as

FNs, and uses machine learning methods for graph-structured data to automatically detect those regions of the brain that are important for localising seizure onsets. To train the GNN, it is sufficient to have seizure onset annotations, but no information regarding the SOZ is needed. We showed the effectiveness of our method in localising the SOZ on real-world data consisting of iEEG recordings from eight human subjects, using two different FC metrics to compute FNs. Our results show a very high accuracy in localising the SOZ. However, we also observed that the GNN exhibits uncertainty in those cases where human analysis was also uncertain, indicating a reliable and safe behaviour to support decision-making.

We believe that this work represents a step towards AI-aided analysis of iEEG data and could potentially lead to faster and more accurate treatment of epilepsy.

References

- Bahdanau D**, Cho K, Bengio Y. Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473. 2014; .
- Bastos AM**, Schoffelen JM. A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Frontiers in Systems Neuroscience*. 2016; 9:175.
- Battaglia PW**, Hamrick JB, Bapst V, Sanchez-Gonzalez A, Zambaldi V, Malinowski M, Tacchetti A, Raposo D, Santoro A, Faulkner R, et al. Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261. 2018; .
- Benjamin O**, Fitzgerald TH, Ashwin P, Tsaneva-Atanasova K, Chowdhury F, Richardson MP, Terry JR. A phenomenological model of seizure initiation suggests network structure may explain seizure frequency in idiopathic generalised epilepsy. *The Journal of Mathematical Neuroscience*. 2012; 2(1):1–30.
- Bianchi FM**, Grattarola D, Livi L, Alippi C. Graph neural networks with convolutional ARMA filters. arXiv preprint arXiv:1901.01343. 2019; .
- Brazier MA**. Spread of seizure discharges in epilepsy: anatomical and electrophysiological considerations. *Experimental Neurology*. 1972; 36(2):263–272.
- Bronstein MM**, Bruna J, LeCun Y, Szlam A, Vandergheynst P. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*. 2017; 34(4):18–42.
- Brown TB**, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, et al. Language models are few-shot learners. arXiv preprint arXiv:2005.14165. 2020; .
- Burns SP**, Santaniello S, Yaffe RB, Jouny CC, Crone NE, Bergey GK, Anderson WS, Sarma SV. Network dynamics of the brain and influence of the epileptic seizure onset zone. *Proceedings of the National Academy of Sciences*. 2014; 111(49):321–330. doi: [10.1073/pnas.1401752111](https://doi.org/10.1073/pnas.1401752111).
- Covert I**, Krishnan B, Najm I, Zhan J, Shore M, Hixson J, Po MJ. Temporal Graph Convolutional Networks for Automatic Seizure Detection. arXiv preprint arXiv:1905.01375. 2019; .
- Defferrard M**, Bresson X, Vandergheynst P. Convolutional neural networks on graphs with fast localized spectral filtering. In: *Advances in Neural Information Processing Systems*; 2016. p. 3844–3852.
- Gadgil S**, Zhao Q, Adeli E, Pfefferbaum A, Sullivan EV, Pohl KM. Spatio-Temporal Graph Convolution for Functional MRI Analysis. arXiv preprint arXiv:2003.10613. 2020; .
- Gersch W**, Goddard G. Epileptic focus location: spectral analysis method. *Science*. 1970; 169(3946):701–702.
- Gilmer J**, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural message passing for quantum chemistry. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70* JMLR. org; 2017. p. 1263–1272.
- Hashiguchi K**, Morioka T, Yoshida F, Miyagi Y, Nagata S, Sakata A, Sasaki T. Correlation between scalp-recorded electroencephalographic and electrocorticographic activities during ictal period. *Seizure*. 2007; 16(3):238–247.
- Jirsa VK**, Stacey WC, Quilichini PP, Ivanov AI, Bernard C. On the nature of seizure dynamics. *Brain*. 2014; 137(8):2210–2230.

- 463 **Khambhati AN**, Davis KA, Lucas TH, Litt B, Bassett DS. Virtual cortical resection reveals push-pull network
464 control preceding seizure evolution. *Neuron*. 2016; 91(5):1170–1182. doi: [10.1016/j.neuron.2016.07.039](https://doi.org/10.1016/j.neuron.2016.07.039).
- 465 **Khambhati AN**, Davis KA, Oommen BS, Chen SH, Lucas TH, Litt B, Bassett DS. Dynamic network drivers of
466 seizure generation, propagation and termination in human neocortical epilepsy. *PLoS Computational Biol-*
467 *ogy*. 2015; 11(12):e1004608. doi: [10.1371/journal.pcbi.1004608](https://doi.org/10.1371/journal.pcbi.1004608).
- 468 **Kwan P**, Brodie MJ. Early identification of refractory epilepsy. *New England Journal of Medicine*. 2000;
469 342(5):314–319.
- 470 **Lachaux JP**, Rodriguez E, Martinerie J, Varela FJ. Measuring phase synchrony in brain signals. *Human Brain*
471 *Mapping*. 1999; 8(4):194–208.
- 472 **Lee HW**, Arora J, Papademetris X, Tokoglu F, Negishi M, Scheinost D, Farooque P, Blumenfeld H, Spencer DD,
473 Constable RT. Altered functional connectivity in seizure onset zones revealed by fMRI intrinsic connectivity.
474 *Neurology*. 2014; 83(24):2269–2277.
- 475 **Litjens G**, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, Van Der Laak JA, Van Ginneken B, Sánchez
476 CI. A survey on deep learning in medical image analysis. *Medical Image Analysis*. 2017; 42:60–88.
- 477 **Lopes MA**, Junges L, Woldman W, Goodfellow M, Terry JR. The Role of Excitability and Network Structure in the
478 Emergence of Focal and Generalized Seizures. *Frontiers in Neurology*. 2020; 11:74.
- 479 **Lopes MA**, Richardson MP, Abela E, Rummel C, Schindler K, Goodfellow M, Terry JR. An optimal strategy for
480 epilepsy surgery: Disruption of the rich-club? *PLoS Computational Biology*. 2017; 13(8):e1005637.
- 481 **Nunez PL**, Srinivasan R, et al. Electric fields of the brain: the neurophysics of EEG. Oxford University Press,
482 USA; 2006.
- 483 **Sanderson M**, Manning CD, Raghavan P, Schütze H. Introduction to Information Retrieval, Cambridge Univer-
484 sity Press. 2008. ISBN-13 978-0-521-86571-5, xxi+ 482 pages. Natural Language Engineering. 2010; 16(1):100–
485 103.
- 486 **Sanz Leon P**, Knock SA, Woodman MM, Domide L, Mersmann J, McIntosh AR, Jirsa V. The Virtual Brain: a
487 simulator of primate brain network dynamics. *Frontiers in Neuroinformatics*. 2013; 7:10.
- 488 **Schindler KA**, Bialonski S, Horstmann MT, Elger CE, Lehnertz K. Evolving functional network properties and
489 synchronizability during human epileptic seizures. *Chaos: An Interdisciplinary Journal of Nonlinear Science*.
490 2008; 18(3):033119.
- 491 **Schlichtkrull M**, Kipf TN, Bloem P, van den Berg R, Titov I, Welling M. Modeling relational data with graph
492 convolutional networks. In: *European Semantic Web Conference* Springer; 2018. p. 593–607.
- 493 **Shah AK**, Mittal S. Invasive electroencephalography monitoring: Indications and presurgical planning. *Annals*
494 *of Indian Academy of Neurology*. 2014; 17(Suppl 1):S89.
- 495 **Simonovsky M**, Komodakis N. Dynamic edgeconditioned filters in convolutional neural networks on graphs.
496 In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; 2017. .
- 497 **Stafstrom CE**, Carmant L. Seizures and epilepsy: An overview for neuroscientists. *Cold Spring Harbor Perspec-*
498 *tives in Medicine*. 2015; 5(6):a022426.
- 499 **Staljanse W**, Strobbe G, Van Holen R, Birot G, Gschwind M, Seeck M, Vandenberghe S, Vulliémot S, van
500 Mierlo P. Seizure onset zone localization from ictal high-density EEG in refractory focal epilepsy. *Brain*
501 *Topography*. 2017; 30(2):257–271.
- 502 **Van Mierlo P**, Papadopoulou M, Carrette E, Boon P, Vandenberghe S, Vonck K, Marinazzo D. Functional brain
503 connectivity from EEG in epilepsy: Seizure prediction and epileptogenic focus localization. *Progress in Neu-*
504 *robiology*. 2014; 121:19–35.
- 505 **Vaswani A**, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I. Attention is all you
506 need. In: *Advances in neural information processing systems*; 2017. p. 5998–6008.
- 507 **Velickovic P**, Cucurull G, Casanova A, Romero A, Lio P, Bengio Y. Graph attention networks. *International*
508 *Conference of Learning Representations (ICLR)*. 2018; .
- 509 **Wang Z**, Yan W, Oates T. Time series classification from scratch with deep neural networks: A strong baseline.
510 In: *2017 International joint conference on neural networks (IJCNN)* IEEE; 2017. p. 1578–1585.

- 511 **Weaver KE**, Chaovalitwongse W, Novotny EJ, Poliakov A, Grabowski Jr TJ, Ojemann JG. Local functional connec-
512 tivity as a pre-surgical tool for seizure focus identification in non-lesion, focal epilepsy. *Frontiers in Neurology*.
513 2013; 4:43.
- 514 **Xu K**, Ba J, Kiros R, Cho K, Courville A, Salakhudinov R, Zemel R, Bengio Y. Show, attend and tell: Neural image
515 caption generation with visual attention. In: *International conference on machine learning*; 2015. p. 2048–2057.
- 516 **Yu B**, Yin H, Zhu Z. Spatio-temporal graph convolutional networks: A deep learning framework for traffic
517 forecasting. *arXiv preprint arXiv:170904875*. 2017; .

Experiments on Simulated Data

Seizure generator from *Benjamin et al. (2012)*

In this experiment we considered a simple network model of seizure initiation presented by *Benjamin et al. (2012)*, and also used by *Lopes et al. (2017, 2020)* to study the effect of network structure on the generation of seizures. The model consists of a network of N bi-stable oscillators

$$\dot{z} = f(z) = (\lambda - 1 + i\omega)z + 2z|z|^2 - z|z|^4 \quad (14)$$

where $z \in \mathbb{C}$. Equation (14) describes a dynamical system with a stable fixed point at the origin of the complex plane (which we consider as interictal), and an oscillating attractor with frequency ω (which we consider as ictal). Parameter λ controls the location of the oscillator in phase space. Nodes are interconnected in a graph described by adjacency matrix $\mathbf{A}$ with a coupling factor β , such that the dynamic of a single node reads:

$$dz_i(t) = (f(z_i) + \beta \sum_{j \neq i} \mathbf{A}_{ji}(z_j - z_i)) + \alpha dW_i(t)$$

where $W_i(t)$ is a stochastic Wiener process rescaled by a factor of α .

All nodes in the model are initialised at the fixed point and, due to the presence of noise and the interaction between nodes, eventually switch to the oscillation state. We identify the activity of the whole system as ictal if any of the nodes meets the condition $|\text{Re}(z_i)| > 1$, and the SOZ as the first node that escapes the fixed regime.

We consider a complete graph without self-loops to describe the interaction of the nodes. The configuration of the parameters is summarised in Table 5. The hyperparameters used for creating the FNs and training the GNN are the same ones that we used for the real iEEG data, and we only report results obtained using PLV as FC metric.

Appendix 0 Table 5. Configuration used for the simulator by *Benjamin et al. (2012)*.

Parameter	Value
N	3
ω	20
λ	0.5
β	0.1
α	0.05

The GNN achieves an almost perfect detection score with a ROC-AUC of 99.61 ± 0.0 and a PR-AUC of 99.69 ± 0.0 (averaged over five runs, evaluated on hold-out test data). Figure 6 compares the generated node activity with the attention scores assigned by the GNN over time. The SOZ channel (green) is assigned the highest attention since the beginning of the seizure until all nodes are simultaneously oscillating, at which point the attention scores converge to be evenly distributed. A similar even distribution is observed in the interictal state, indicating that the network has correctly learned to identify the SOZ electrode without defaulting to assign a high score to just one electrode. This behaviour is confirmed by the spikes in attention assigned to channels 0 and 1, which happen as soon as the node dynamics escape the fixed-point attractor.

The Virtual Brain Simulator

In this experiment we use The Virtual Brain simulator (TVB) (*Sanz Leon et al., 2013*) to model a patient with temporal lobe epilepsy.

We follow the same approach described in TVB's documentation to configure the simulator.¹ We assign the Epileptor neural mass model (*Jirsa et al., 2014*) to all the controllable brain regions of

¹https://github.com/the-virtual-brain/tvb-root/blob/master/tvb_documentation

Appendix 0 Figure 6. Top: a clip showing the generated activity of a 3-node simulator, compared to the attention coefficient assigned by the GNN at each node over time. Colors indicate the same node in both plots.

TVB. We set the epileptogenicity of the right limbic areas (rHC, rPHC and rAMYG) to -1.6 , the superior temporal cortex (rTCI) and the ventral temporal cortex (rTCV) to -1.8 , while for all other areas to -2.2 . The remaining parameters are kept as default. The hyperparameters used for creating the FNs and training the GNN are the same ones that we used for the real iEEG data.

We select a subset of 34 sEEG virtual sensors among the ones provided for the default subject of TVB. Of this subset, electrode 33 shows strong epileptogenic activity, while electrodes 18, 19, and 20 show mild activity. We generate clips of roughly 1 minute at 20Hz so that there is a simulated onset in the middle of each clip. An example of a generated clip is shown in Figure 7.

The GNN achieved an average detection ROC-AUC of 98.87 ± 0.18 and an average PR-AUC of 99.18 ± 0.07 (averaged over five runs, evaluated on hold-out test data). The electrode with a strong ictal activity is consistently assigned a maximum score of 1 by all models and electrode 19 is also ranked in the top-5 electrodes (see Figure 8).

GNN training details

We consider each patient separately and train a GNN from scratch to build patient-specific models. The GNN architecture is the one given in Equation (11). The ECC layer has 32 output units with ReLU activation and a kernel-generating network $f(\cdot)$ consisting of a two-layer MLP with 32 hidden units and ReLU activation. All parameters of the layer are regularised with an L_2 penalty with a factor of 10^{-5} .

The MLP classifier following the Attn-RO readout has 2 layers, with the hidden one having 32 units and ReLU activation and with 25% dropout in-between. Both layers are regularised with an L_2 penalty with factor 10^{-5} .

The model is trained using Adam, with a learning rate of 10^{-3} and a batch size of 32 graphs. The model is trained to convergence with 10 epochs of patience, using the data from $[0.2 \cdot n]$ seizures selected randomly (n being the overall number of available seizures) for early stopping.

Appendix 0 Figure 7. A virtual seizure generated with TVB. The vertical line denotes the annotated seizure onset in time.

Appendix 0 Figure 8. Top-10 electrodes with averaged rankings. Bold labels indicate that the corresponding electrode showed ictal activity. As desired, electrode 33 shows strong epileptogenic activity.

Baseline training details

The baseline is a simple 1D convolutional neural network (CNN) based on the architecture described by *Wang et al. (2017)*. The CNN operates directly on iEEG time series and hence does not take into account any graph-based representation for the data. Similarly to how we create the input-output pairs for the GNN, here we consider windows of size T taken at a stride of k/f_s for the interictal class and stride $1/f_s$ for the ictal class, and we associate to each window a class label corresponding to the majority class of $y(t)$ in the corresponding window.

In particular, we shrink the model to make it comparable in terms of number of parameters and depth to the GNN one, and also to prevent overfitting (which we experimentally encountered as a significant problem with the model). We consider a single convolutional layer with a kernel of size 3, 8 output channels, and ReLU activations, followed by a global average pooling and a single-layer MLP to output the classification decision. We train the model using Adam with learning rate 0.001, batch size of 32 and early stopping with a patience of 5 epochs.