Volume 0 (1981), Number 0 pp. 1–16 COMPUTER GRAPHICS *forum*

Neurosurgery Network Pattern Analysis with 2nd Generation Ensembles

N. Nipu, S. Zhao, B. Maharathi, J. A. Loeb, and G. E. Marai

University of Illinois Chicago, Chicago, Illinois, USA
mnipu2, szhao69, bmahar2, jaloeb, gmarai@uic.edu

Abstract

Evolving measuring and computing capabilities, along with increasingly complex problems or models, are resulting in a new type of dataset: second-generation ensemble data. Like first-generation ensemble data, these data consist of large-scale, repeated measurements of the same process or phenomenon, and often have a spatial component. Unlike older datasets, they typically require the extraction and aggregation of novel complex features, which may be generated through direct measurements rather than simulations, and appear in a wider range of application domains, including neuroscience. We describe an interactive visual analysis solution for this type of second-generation ensemble data, related to the study and planning of surgical interventions in epilepsy treatment. As part of this solution, we introduce a dynamic community abstraction in conjunction with analysis algorithms for feature extraction and aggregation, registration techniques to correlate and project sample data, and custom visual encodings to support the analysis of conserved network patterns. A quantitative and qualitative evaluation with domain experts at four sites demonstrates the effectiveness of this solution. We discuss this approach and solution in the context of second-generation ensemble data analysis, along with the challenges of working with this type of data.

Keywords: Life Sciences, Health, Medicine, Biology, Bioinformatics, Genomics; Feature Detection, Extraction, Tracking & Transformation; Dynamic Community Extraction; Temporal Data; Ensemble Visualization

1. Introduction

Some epilepsy patients do not respond to medications [E*12, JC15], but may respond to surgical intervention. Surgeons cannot replicate on-demand epileptic seizures in the clinic as they seek to plan interventions. However, neurologists can instrument the patient's scalp with electrodes, then measure and observe electrical activities generated repeatedly by electroencephalography (EEG) recordings [AE*18], in the hope of identifying repeated propagation patterns between affected regions of the brain [M*19]. Clinically, in such cases, the presence of conserved patterns is used to determine whether the patient is a suitable candidate for surgery: consistent patterns indicate an underlying physiological structure that can be addressed surgically. If repeated spatial propagation patterns exist, the neurologists then seek to identify the origin region of the propagation pattern, and then communicate the findings to the surgeons. Surgically rescinding the origin region of such patterns has the potential to disrupt future seizures [dA01].

The example above is a modern instance of an **ensemble dataset** and problem. Ensemble data, introduced as a data concept in 1992 and in the data visualization literature in 2005, are collections of time-varying datasets (ensemble members) that model the same object or process. Each ensemble member is characterized by different input conditions, models, or parameters containing multivariate attributes. First-generation ensemble datasets have been commonly generated through high-performance computing simulations in engineering and climate modeling [WHLS18]. They benefit from visual and statistical analysis to identify patterns, trends, and anomalies. The main visualization challenges are due to the large scale of the data, which leads to visual cluttering, slow rendering times, and a significant learning curve for experts who have a low level of data visualization literacy.

Advancements in measurement and computing technologies are now producing spatial ensemble datasets, where the challenges are no longer visualizing many members, but identifying the structures and features that must be extracted before meaningful comparison. In several recent visualization works [F*19, KBL22, v*23, N*25], including contrail detection [N*23] and viscous-finger analysis [L*18], the features of interest were not explicitly available and thus required domain-guided abstraction, feature extraction techniques, aggregation, and summarization. Motivated by these observations, in this paper, we introduce the term **2nd generation ensemble datasets** to describe cases where relevant spatial abstractions must be defined or derived prior to visual, statistical, or machine-learning analysis. Examples of such applications include viscous-finger structures in fluid mechanics and clinically measured brain activity, where meaningful spatial patterns ("hard to define, but recognizable when seen") require extraction before com-

parison. Unlike first-generation ensembles, where features such as temperature, pressure, or wind speed are explicitly available, these datasets require determining what should be measured or aggregated before analysis can be conducted. This distinction helps clarify why recent machine learning surrogate approaches, which assume pre-defined output fields or labels, are not directly applicable until the necessary extraction has taken place. Second-generation ensemble datasets can also be generated through direct measurements rather than simulations [Z*17, C*20, W*19, W*24]. As such, they may contain fewer members but demand a more complex spatial transformation. In our epilepsy study, clinically meaningful features such as "spike propagation pattern" or "seizure onset region" are not explicitly labeled and must be inferred through spatio-temporal network analysis, dynamic community detection, and projection techniques. We therefore employ the term "second-generation ensembles" in this context to position our work among others that share similar extraction-first requirements.

This work makes a two-fold contribution. First, we present a visual computing framework for analyzing spatial pattern conservancy in dynamic brain epilepsy data. Our framework constructs a visual computing solution to extract meaningful spatial measures of pattern conservancy across multiple ensemble members and pinpoint the specific spatial locations that may be responsible for seizure onset. From a visual computing perspective, in addition to feature extraction, our solution introduces a social network-inspired dynamic network abstraction and registration algorithms to cross-correlate the spatial data and to map the data to 2D for analysis. On the front-end, a multi-view interface leverages custom encodings to support visual analysis of the data and model. We implement the resulting approach into an interactive system and evaluate it with domain experts. Second, informed by this application and by related visualization works, we introduce the concept of 2nd-generation ensemble data. We reflect on how such datasets relate to a broader family of ensemble problems, highlighting the necessary shift in visual computing approaches from statistical summarization to more sophisticated data abstraction and feature extraction techniques. Last, we discuss the characteristics and challenges of working with this type of data.

2. Related Work

Ensemble visualization. Visualization techniques have been proposed to analyze ensemble data in various domains such as meteorology and climatology [P*09, S*10], physics [B*15, HHB16], oceanography [H*14b, H*13], and biology [DPJP20, HSSK16]. Statistical summaries, such as mean, variance [MWK14, WMK13], probability distributions [PWH11], and clustering methods [FBW16, FKRW16], are often employed to simplify ensemble data, and statistical visual encodings and coordinated multiple views [N*23, L*18] are typically leveraged. For a broader overview, readers are referred to recent surveys and STAR reports on ensemble visualization [WHLS18, KH13].

Since 2017, ensemble data visualization emphasizes innovative feature extraction across spatial domains, including computational fluid dynamics simulations [N*23, L*18] and climate simulations [F*19, KBL22]. Other work analyzes directly measured ensemble data in medicine [Z*17, C*20], threat assessment [X*19a], or music [YC23], again with a focus on innovative feature extraction. Our work falls into this second category and also deals with directly measured ensemble data. Our work is also the first to examine these new types of ensemble datasets and problems, which we term "2nd generation ensembles", in the context of visualization research.

Spatiotemporal visualization. Visualization of multivariate time-dependent spatial data appears in domains such as medicine, climate research, or engineering [KH13]. Different visualization techniques such as glyphs [N*23, TSWS05], time histograms [KBH04, AFM06], color and texture [HE99], layering [RBG07, KML99], flow maps [AA11], volume rendering [N*23, CS99] and various interaction features such as brushing [HS04, K*08], focus+context [M*08] are used to display spatiotemporal attributes of the data. For an extended discussion, readers may consult recent surveys and STAR reports on spatiotemporal visualization [DMK05, A*08, AA06]. Our system employs a combination of aggregation, 2D projection, and clustering techniques to visualize spatiotemporal brain electrode data. This approach is further enhanced by various visualization techniques, including volume rendering, color, line curves, and interaction techniques such as brushing and focus+context, in addition to the line charts for EEG data preferred by neuroscientists.

Brain network visualization. In brain network visualization, common visual representations are matrix-based representations [M*17, SA*10, A*13, SA*10, BPF14], node-link diagrams [S*22, WCLE05, S*05, X*19b], or a combination in a hybrid approach [M*15, Y*17]. Several systems look at functional connectivity [R*15, ZFB10, HHK12], including EEG and Magnetoencephalography (MEG) recordings with multiple user perspectives [XWH13, SSWZ23], 2D and 3D network views [H*15, H*11], or projections on three anatomical planes to represent human brain networks [WCLE05]. These works focus on single-session and do not allow comparison across multi-sessions. Fujiwara et al. [F*17] proposed an interactive visual analytics tool for analyzing resting-state functional connectivity data across multi-sessions. However, they do not seek to extract repeated patterns. Our work leverages similar visual encodings but is focused on pattern extraction. Community detection algorithms have been utilized in brain connectomics to analyze brain networks, primarily to identify functionally co-active regions within a single recording session [M*17], as well as to identify modular structures in static functional brain networks derived from a single-session [GZYL23, NBB17]. These approaches are typically applied to networks with fixed nodes that represent stable brain regions, and do not account for variability across multiple sessions. In contrast, our work focuses on epilepsy ensemble data, characterized by repeated measurements and variability in electrode activation patterns across multiple sessions. We aim to extract spatially and temporally conserved activation patterns across multiple sessions, enabling the identification of consistent activated regions over time.

3. Methods

This section describes our data, interdisciplinary collaboration settings, analytical requirements, and computational methods for feature extraction, community detection, and pattern conservation.

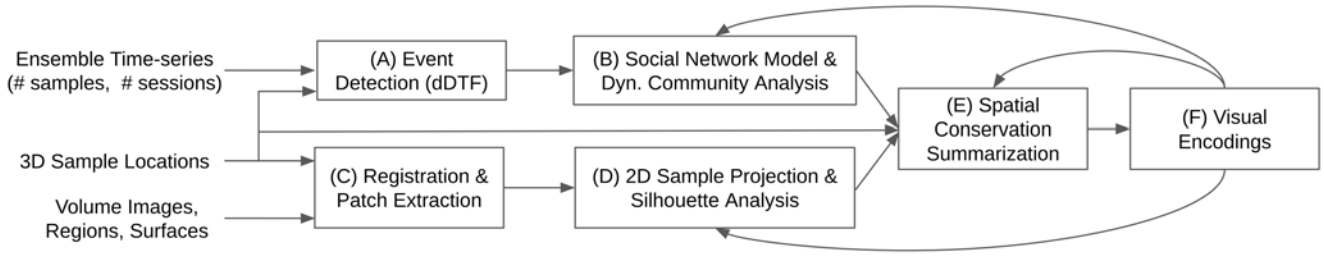

Figure 1: Feature extraction pipeline for network pattern analysis. The front-end interface helps communicate results as well as steer the model development process. Input data are processed through (A) Event Detection (dDTF) to support the (B) Social Network Model & Dynamic Community Analysis, and through (C) Registration & Patch Extraction, followed by (D) 2D Sample Projection & Silhouette Analysis. The resulting outputs are combined to extract (E) Spatial Conservation Summarization, and the results are visualized through (F) Visual Encodings.

3.1. Background Problem and Data

Epilepsy is a common neurological disorder characterized by recurrent seizures [OSR12, KSCW18]. Seizures manifest as dynamic, spatial networks of electrical spikes that span multiple brain regions, and which sometimes correlate with the location of brain lesions. Although seizures are not frequent events, abnormal spike networks are not only confined to seizures but also occur, in fact, frequently between seizures. These abnormal spatial spike networks can be studied in the clinic by instrumenting the patient's scalp with electrodes (70 to 120 electrodes) and tracking the electrical activity at these electrode samples across a handful of recorded sessions (2 to 5 sessions) for brief intervals (120 to 600 seconds). The electrodes are typically organized in patches, which are then placed so that they cover specific anatomical regions of the brain. Volume images (Computed Tomography (CT)/Magnetic Resonance Imaging (MRI)) of the patient's brain are also acquired in the clinic, allowing the identification of regions of interest and reconstruction of the brain surface and brain lesions. Regions are physician-marked areas of the brain, while patches are groups of electrodes placed on the surface. A patch may cover multiple regions. These records form a large spatiotemporal ensemble.

Observing raw EEG signals allows for detecting spike events and analyzing the propagation of electrical activities. Neurologists typically review 10 to 20 EEG signals at a time [PP14] in 10 to 20-second intervals. Identifying potential seizure time windows involves the definition of interictal spikes [M*19]. Interictal (or interseizure) spikes refer to brief electrical signals lasting 250 ms or less in the brain. These signals consist of a short, sharp wave followed by a slow wave that lasts longer [dA01]. By leveraging this definition, one could better evaluate potential surgery candidates who have not responded to medication.

3.2. Project Setting and Computational Modeling

Our solution was designed through a multi-disciplinary remote collaboration among several research and clinician groups at 4 sites over 2 years. Our visual computing team met monthly, in addition to smaller weekly meetings, with a team of neurologists and bioengineers at our institution to develop the methodology for feature extraction and pattern conservancy analysis. Our two teams also met jointly with surgeons at 4 clinical sites to coordinate data collection, present results, and collect feedback.

The multi-disciplinary approach is necessary due to difficulties in collecting, interpreting, and analyzing the neurological spike data. Interpretation and analysis involve neurologists and surgeons working jointly. Together, they visually examine the temporal EEG data from the electrodes in conjunction with medical imaging data across multiple tools, then use domain knowledge to identify regions of interest. In turn, bioengineers and visual computing specialists create computational methods to post-process the EEG data, detect spike events, construct activation networks based on electrode spikes and propagation of the spikes to other electrodes, and visualize results in 2D and 3D space.

Specifically, our bioengineering collaborators use the direct-directed transfer function (dDTF), an effective connectivity measure, to detect spike events in EEG data and map the network behavior of the spikes within a specific region in the human epileptic brain [M*19]. The dDTF method, originally published in the '80s and with a basic algorithm available under Matlab, is well-vetted and well-established, and has been adopted by many research groups. Each event is characterized by the spiked electrodes' spatial location, the activation time of the spiked electrodes, a list of activated/spiked electrodes, and their propagation to other electrodes. The result is, however, a set of extremely dense networks where dense connections are then filtered based on node activation thresholds. From this dataset, one can further calculate and report for further analysis the number of spikes recorded by each electrode node and the number of outgoing and incoming edges for that node, as yielded by the dDTF propagation data for each node.

The sequence and onset of the spikes are then painstakingly analyzed across several tools, including Matlab and python, to gain insights into the patterns of epileptic activity in the brain. Our bioengineering and neurology collaborators had found in previous analyses that spike onset regions in the brain, as identified by surgeons, correlate with those brain regions that have the greatest number of outgoing spike propagations above a defined maximum threshold within a selected grid. Counterintuitively, the spike onset regions did not correlate with those brain regions with the highest spikes. At the same time, the spike spread was associated with those

brain regions with the highest number of incoming spike propagations. Although these analyses had only been performed at the level of a single patch of electrodes. Our collaborators expected the same to hold true in multiple patch analysis. The approach, however does not scale well, by itself, to multiple patch analysis, nor does it support the analysis of pattern conservation across multiple sessions.

3.3. Requirements Analysis

We summarize the functional requirements for the project as tasks below [Mar18], where tasks T3-T5 reflect the project needs in terms of model development, and T1, T2, T6, T7 reflect primarily the needs of the surgeons:

- T1. Identify the spatiotemporal location of regions of high electrical activity (i.e., spike locations).
- T2. Identify the spatial source (onset) and spread of high electrical activity (i.e., spikes).
- T3. Explore the electrode EEGs in detail to verify the output of algorithms.
- T4. Analyze the spatial electrical activity within a regional patch at different electrical activity thresholds.
- T5. Analyze the spatial electrical activity among regional patches at different activity thresholds.
- T6. Analyze whether spike onset locations correlate with lesion locations (if any).
- T7. Determine whether the spike onset and spread (i.e., network or pattern) are conserved across multiple sessions for the same patient.

Non-functional requirements included visual scaffolding that would include familiar spatial and non-spatial visualizations [Mar15], efficient scalability to large brain networks and large EEG data, and online availability for remote collaborators.

3.4. Abstraction, Feature Detection and Extraction

3.4.1. Dense Social Network Abstraction

Our bioengineering collaborators had previously attempted and discarded standard graph modeling and graph analysis algorithms such as spectral clustering or centrality measures. These approaches did not work well, given the highly edge-dense characteristics of the propagation networks and their dynamic nature. To circumvent these issues, we abstract and model the propagation network as a dense social network instead, while continuing to use the dDTF event detection algorithm. In this analogy, each electrode is considered a node (similar to an individual in a social network), and the propagation between electrodes is modeled as an interaction or edge. Community analysis algorithms have been introduced in the study of social networks and biological networks [GN02, H*14a] to help identify similar groups of nodes based on their network activity, similar to the problem we explore. In this type of network analysis, a community refers to a group of nodes with a high level of interconnectivity with each other and sparsely connected to the remainder of the network [YLL10]. By adopting this abstraction, we would be able to pre-filter highly interconnected electrodes (T4), and avoid computational scalability and visual clutter challenges associated with our ensemble problem. However, to operate this abstraction, we need to first define and characterize spike patterns. Figure 1 shows the abstraction and feature extraction components necessary to support pattern conservation analysis.

3.4.2. Characterizing Spike Patterns

Based on bioengineering and neurologist collaborator input (Sec. 3.2), we characterize the spike patterns in terms of the spatial location of the electrodes that registered electrical activity spikes (Fig. 1.A). We also use measures related to the total number of spikes registered at each electrode and their propagation pattern. Accordingly, we characterize the spike patterns as follows, using the following three entities for each electrode, in addition to spatial location:

- **Spike frequency** as the total number of spike occurrences per session at that electrode.
- **Spike onset** as the total number of outgoing propagation edges at that electrode.
- **Spike spread** as the total number of incoming propagation edges at that electrode.

Using these data, we then seek to detect which electrodes consistently act together as a community across sessions, where "act" denotes an electrode registering high electrical activity, i.e., spikes. This analysis has the potential to implicitly capture the conservation of spike patterns across sessions (Fig. 2). Specifically, spike propagation detection captures the temporal ordering and direction of electrical events, whereas community detection identifies spatially conserved co-activated electrode groups across sessions.

3.4.3. Electrode Community Analysis

Because our goal is to detect communities that are stable across sessions, we adopt a temporal trade-off community detection technique based on the Louvain algorithm [BGLL08], where the community defined at time step t depends on the topology of the network and communities at time step $t-1$ [AG10]. This approach helps detect stable communities, as modifications in the communities reflect actual modification of structural changes in the network rather than simply artifacts of the algorithm [RC18]. We use the electrodes as the social network nodes and the propagations between electrodes as edges. To extract community conservation across recorded sessions (T3), we first aggregate the network connections to generate a weighted network for each session, where weight corresponds to the number of propagations between two electrodes. We then apply our modified version of the Louvain method to identify communities of electrodes conserved via consistent activations across the sessions (Fig. 1.B).

The Louvain method is a greedy algorithm based on the modularity gain M of a partition [New06], which measures network density inside communities compared to links between communities [BGLL08]:

$$M = \frac{1}{2w}\sum_{x=1}^{N}\sum_{y=1}^{N}\left(E_{xy} - \frac{e_x e_y}{2w}\right)\delta(C_x, C_y) \tag{1}$$

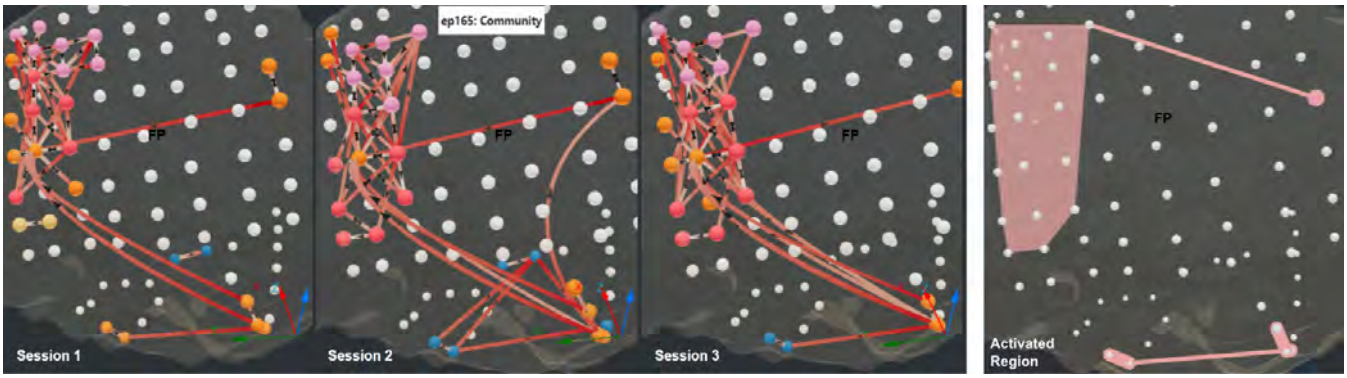

Figure 2: Electrode community analysis of patient ep165, showing the result of the community detection algorithm across three sessions. A small group of electrodes inside the region is responsible for the high electrical activity across all three sessions. These highly connected electrodes are grouped into three communities, shown in blue, orange, and pink. Activated region indicating that not all community electrodes are conserved, but the overall pattern is conserved across sessions. Note that the connectors in the activated region view represent geometric nearest-neighbor links added solely to illustrate spatial continuity; they do not correspond to the actual propagation paths observed in individual sessions.

where N is the total number of electrodes, w is the total number of propagations, E_{xy} is the weight, $e_x = \sum_x E_{xy}$ and $e_y = \sum_y E_{xy}$ are the sum of the propagation weights attached to electrode x and y, C_x and C_y are the communities to which electrode x and y belong and $\delta(C_x, C_y)$ defines whether electrode x and y are in the same community:

$$\delta(C_x, C_y) = \begin{cases} 1, & \text{if } x \text{ and } y \text{ are in the same community} \\ 0, & \text{Otherwise} \end{cases} \tag{2}$$

The modularity formula indicates that modularity gains focus solely on the inner part of each community. The modularity change will only affect the communities if an electrode is moved. Hence, during the iteration process, it is necessary to calculate the modularity increment and difference solely without considering the overall modularity before and after relocation [T*21]. The formula for calculating the modularity gain M_c by moving an electrode x into a particular community c is as follows:

$$\begin{aligned} M_c = & \left(\frac{\sum P_{in} + 2e_{x,P_{in}}}{2w} - \left(\frac{\sum P_{tot} + e_x}{2w} \right)^2 \right) \\ & - \left(\frac{\sum P_{in}}{2w} - \left(\frac{\sum P_{tot}}{2w} \right)^2 - \left(\frac{e_i}{2w} \right)^2 \right) \end{aligned} \tag{3}$$

where $\sum P_{in}$ is the sum of the total propagations of all electrodes in the community, $\sum P_{tot}$ is the total number of propagations of all electrodes connected to the community, $e_{x,P_{in}}$ is the total number of propagations of electrodes connecting to electrode x within the community.

In short, the Louvain algorithm begins by allocating each electrode of a session to a separate community. It then proceeds to iterate over each electrode within the inner loop. For each electrode, the algorithm computes two measurements– firstly, it calculates the modularity gain when the electrode is incorporated into the community of any of its neighbors; secondly, it selects the neighbor that provides the highest gain and assigns the electrode to that corresponding community. This process continues until no further movement leads to an increase in the gain.

In our modified algorithm to detect dynamic communities, at timestep t, instead of assigning each electrode to an individual community as the initial step, the algorithm begins by grouping the electrodes into communities identified at timestep $t-1$ and then continues with the remaining steps of the Louvain algorithm to detect communities. In agreement with the model development requirement of threshold-based analysis, we repeatedly apply the dynamic community detection algorithm offline to identify communities over different percentiles. The percentile thresholds were determined in consultation with clinical collaborators to emphasize highly propagated, clinically relevant electrodes, as visualizing the full unfiltered network obscures meaningful activity. Generally, different percentiles change the number of communities yielded. However, highly activated electrodes are conserved in all cases, and high thresholds yield a reasonably low number of communities (Fig. 2). The algorithm produces stable electrode communities across multiple patient sessions (Sec. 5).

3.4.4. Electrode Activation Pattern Conservation Summarization

We propose an algorithm (see Supplemental Materials) to extract activation patterns that are conserved across repeated recorded sessions, even though the exact propagations may differ between sessions (Fig. 1.E). To extract the conserved pattern, we first filter based on the quantile threshold and identify each sub-network, which consists of a list of connected components per sample (Fig. 3.A). Next, we apply a customized Breadth-First Search (BFS) over the 2D anatomical grid for each region to identify spatially coherent activated electrode groups within each subnetwork. For every activated electrode, BFS explores 4-way neighbors to

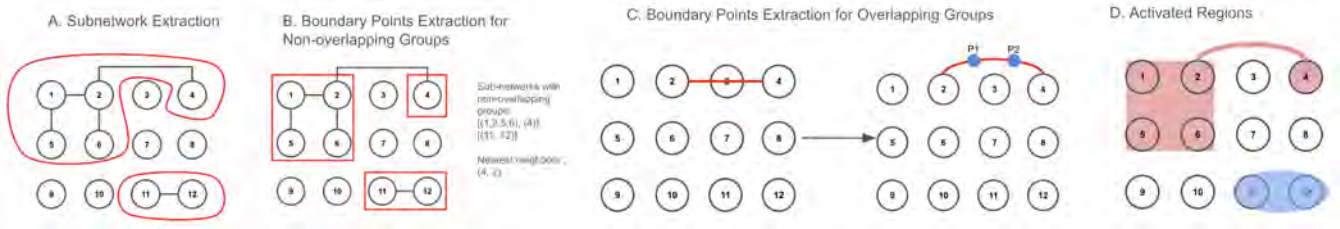

Figure 3: Electrode conserved pattern extraction: (A) For a selected session, sub-networks corresponding to each brain region are identified; (B) Within each sub-network, non-overlapping groups are identified, and their boundary points are extracted; (C) For overlapping regions, boundary points are extracted by using two nearest electrodes; (D) Extracted conserved pattern.

form spatial clusters. To prevent visual or electrode overlap, we explicitly track electrode assignments. Once an electrode is assigned to a group, it is marked as visited and excluded from all subsequent BFS traversals. This process ensures that electrodes from different regions are not incorrectly grouped and that each group only includes activated electrodes without overlaps. After grouping, we extract a distinct set of 3D coordinates that are at the edges of these groups (i.e., boundary points) (Fig. 3.B). To connect different groups of electrodes in the same sub-network while considering overlaps, we check pairs of electrode groups to identify the electrode pairs that are closest. We then draw a line between these electrodes and check for any electrodes that lie along that line. If any are found, we place points before and after the overlapping electrode with a slight variation to prevent the line from overlapping. This process is repeated until no overlaps remain, and we store these points for the connected components (Fig. 3.C). Finally, we display a convex mesh for groups of four or more points, representing a single electrode with a sphere, and using lines for electrode pairs, along with the connected components (Fig. 3.D).

3.4.5. Electrode Patch Extraction

Because the ensemble data does not specify the electrode assignment to patches, we leverage K-Means clustering to automatically group electrodes and extract patches based on the 3D electrode coordinates (Fig. 1.C). K-Means was chosen because it is a well-established and widely used clustering algorithm. The goal of this method is exploratory, where we aim to determine whether an automated approach can effectively replicate the manual grouping of brain regions performed by clinicians. In this context, K-Means offers a straightforward and reproducible alternative that yields spatially coherent electrode clusters for expert validation. Although we explored other clustering methods, they did not yield better results. We identify the number of patches (K) via the silhouette score [Rou87] with varying cluster numbers (2 to 25), an initial seed value of 0, and a max of 300 iterations, which yielded the best result. Finally, we apply the K-means algorithm with the determined parameters to calculate the electrode assignment to patches.

3.4.6. Electrode 2D Snapshot Extraction

3D brain views allow for observing the spatial electrode distribution across the scalp, facilitating the spatial analysis of activation patterns and propagation trends. Despite this advantage, foreshortening and occlusion effects make it difficult to analyze propagations that involve multiple brain regions simultaneously. We have developed an Electrode 2D Snapshot Extraction algorithm to address this issue. This approach maps the 3D data to 2D while preserving detailed information about electrode positions and activities (Fig. 1.D) (T2, T4).

Several tools have been developed to facilitate the visualization of electrode placements and their activity within the brain [G*13, HVHM21]. However, a notable limitation is their integration within specific software environments, such as MATLAB or Python, as opposed to web-based applications.

Our approach provides insight into each patch's spatial arrangement and orientation to ensure physicians can accurately interpret the spike data. This insight helps in complex cases where a single patch spans multiple brain regions. We employ a Multi-view Orthographic Projection [CP78], which accurately identifies structured layouts. While 3D-to-2D projection can theoretically introduce spatial distortion, this method aligns with established clinical workflows, where experts routinely analyze electrode layouts and propagation patterns in 2D grids. The orthographic projection maintains electrode proximity and orientation, and domain experts verified that the resulting 2D representations remained anatomically interpretable and consistent with their clinical expectations. This projection preserves the spatial information of electrode placements and supports data dimensionality reduction. For instance, projecting electrodes from the lower temporal lobe (TL) along the vertical axis (i.e., eliminating the Z-axis coordinates) yields good 2D patch arrangement and orientation.

Next, we employ standard deviation analysis for electrode layout orientation to optimize electrode positioning with maximal spatial accuracy post-dimensionality reduction.We also calculate electrode Euclidean distances, utilizing Interquartile Range (IQR) to adjust electrode thresholds dynamically to distinguish electrode rows and columns. For each brain region and view, we first compute the Euclidean distance between consecutive electrodes in the projected plane. We then compute the first and third quartiles, Q1 and Q3, and the interquartile range $IQR = Q_3 - Q_1$. Distances are then filtered using Tukey's rule to remove outliers:

$$d_i \in [Q_1 - 1.5\,\mathrm{IQR},\ Q_3 + 1.5\,\mathrm{IQR}]. \tag{4}$$

We then compute the mean of these inlier distances, denoted μ_{inlier}, and define the adaptive grouping threshold: $\tau = \mu_{\text{inlier}} + \delta.$, where δ is a small view-specific offset (e.g., 5 for the left–right view, 8 for the top–bottom view) to allow for natural spacing variations.

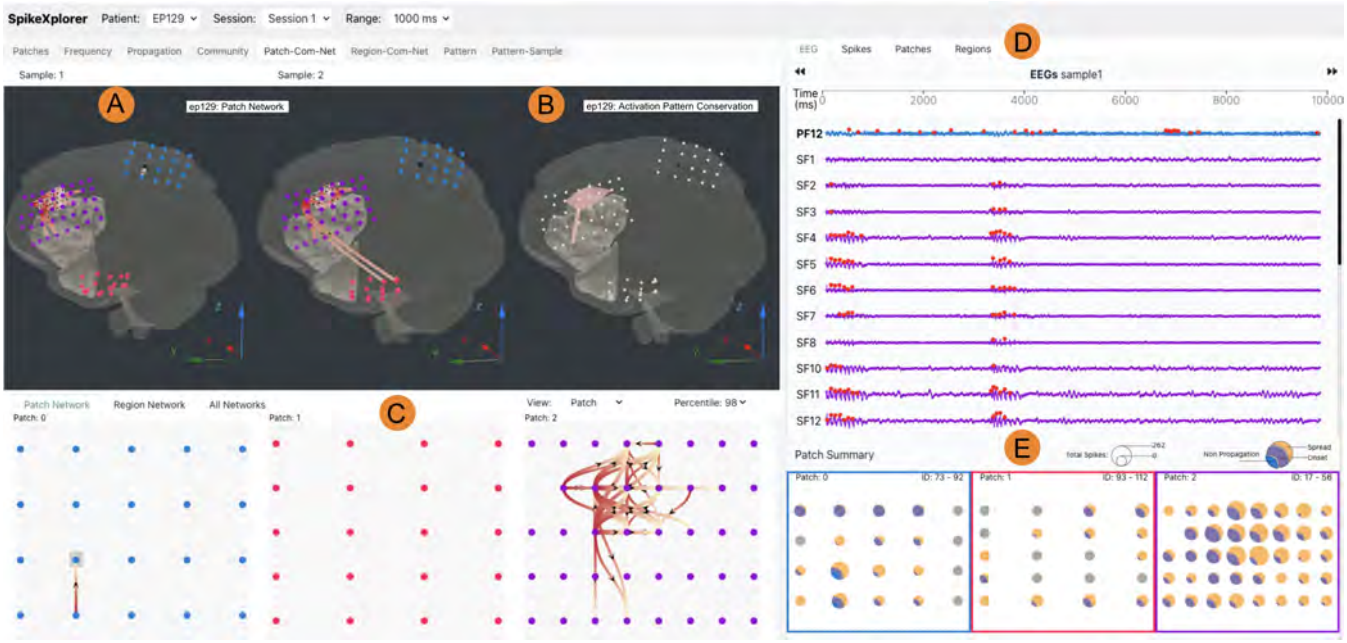

Figure 4: Analyzing conserved electrical activity patterns in epilepsy patients based on ensemble spatial measurements. (A) 3D view for the brain volume, electrodes, lesions, and activity patterns extracted through a social network abstraction, in two recorded sessions; most propagations occur within a patch. (B) Pattern conservation view across the two sessions, calculated through a dynamic community analysis approach. (C) Patch network view showing filtered electrode propagations within-patch; high electrical activities are visible in patch #2. (D) EEG panel view showing EEG line charts and spike occurrences, confirming high spike occurrences in patch #2. (E) Patch summary view showing, with size, the total number of spikes registered, respectively, with color no activation (grey), no propagation (blue), onset (purple), and spread (yellow) spike attributes. Most high-spiking electrodes have a high spike-spread but a low spike-onset.

Whenever the distance between two consecutive electrodes exceeds τ, we start a new group. This procedure allows the threshold τ to adapt to the actual spacing structure of each electrode array while remaining robust to occasional outlier distances.

A second procedure addresses irregularities in the resulting matrix alignment (e.g., in patches where one of the rows contains a different number of electrodes than other rows) as follows:

Procedure 1 Irregular Alignment

```
function AlignMatrix(matrix, orientation)
    max_row_length ← max(len(row) for row in matrix)
    longest_row ← max(matrix, key = len)
    aligned_matrix ← None
    for i, row in enumerate(matrix) do
        for j, electrode in enumerate(row) do
            if len(row) == max_row_length then
                aligned_matrix[i][j] ← electrode
            else
                for each e in longest_row do
                    Calculate the dists between electrode and e
                end for
                min_dist_idx ← dists.index(min(dists))
                aligned_matrix[i][min_dist_idx] ← electrode
            end if
        end for
    end for
    return aligned_matrix
end function
```

This process yields 2D layouts across different brain regions, while preserving the size and orientation of patches (see Supplemental Materials). We verified the output on 23 patches from five patients (four with three patches, one with 11). Most patches followed the grid orientation correctly, with one exception that was manually adjusted. Additionally, there were two instances where a lone electrode was added to the larger grid and which we processed manually.

4. Visual Computing Front-End

In this section, we introduce the visual interface developed to support analytical tasks, detailing the coordinated multi-view design and key interaction components. Medical applications tend to present few opportunities for visual encoding innovation, as visual literacy among clinician clients is typically low. Nevertheless, in this work, we had success with scaffolding from encodings pre-existing in the application domain, with a blend of custom encodings to serve a novel application (Fig. 1.F). Adoption of the encodings was further supported through innovative and customized algorithms to support activities critical to the application domain.

We followed a parallel prototyping approach [D*11] due to its success in eliciting more thorough and helpful feedback than serial

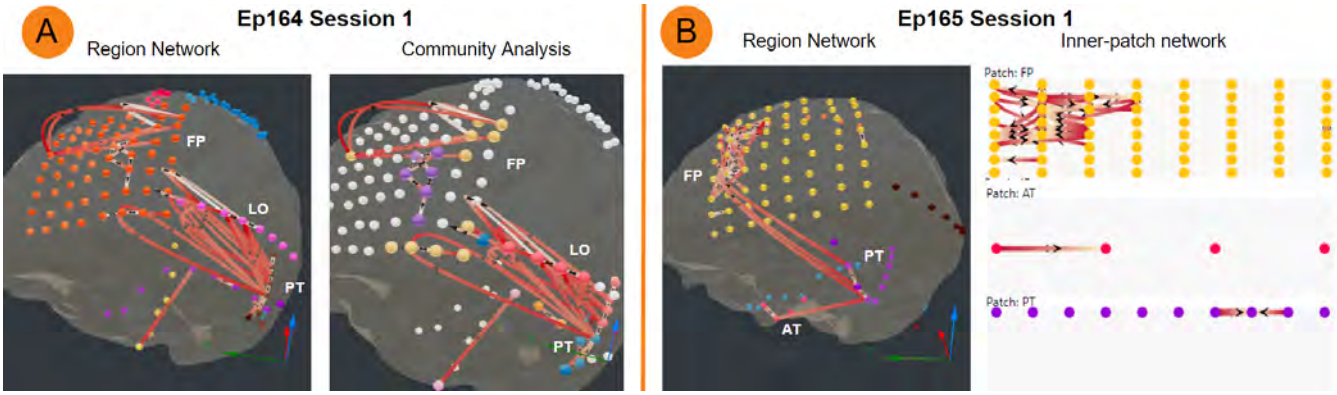

Figure 5: Multi-region spike propagation and activation analysis. Session 1, with spike propagation across regions above the 99th percentile, is selected for ep164 and ep165. (A) ep164 region network showing the propagations between the frontal-parietal (FP), lateral occipital (LO), and posterior subtemporal (PT) regions. The community detection algorithm detected four community groups. (B) Region network and inner-patch network reveal ep165 has high spike activity in the frontal-parietal (FP) area and secondary electrical activity in the posterior subtemporal (PT) and anterior subtemporal (AT) regions.

prototyping. The final design leverages qualitative feedback from our collaborators. The top design of our front-end is based on four coordinated main views supporting the tasks we identified earlier: 1) Electrode Propagation and 3D Brain Network assists in providing the spatio-temporal information of the patient brain, lesions, electrodes, and their activation pattern over time, electrode community and patch network activities across multiple sessions (Fig 4.A, B); 2) Brain 2D Network Exploration supports the examination of the electrode propagations based on patches and regions (Fig 4.C); 3) The EEG Panel allows exploring the EEG propagation for multiple electrodes (Fig 4.D); 4) Patch Summary shows the summary of each electrode grouped by patches and regions (Fig 4.E). The views are linked through color, as well as through brushing and linking.

4.1. Electrode Propagation and 3D Brain Network View.

This view helps identify the spatial location of the brain lesions and electrodes based on the patient-specific brain model in a 3D space (T1, T6). Additionally, this view aims to help specify the electrode spike occurrences, their location over time, and network propagation for each session (T2, T4, T5, T7). Due to the complexity of the necessary information and the visual scaffolding requirement, the view is organized in multiple tabs according to the task analysis and in side-by-side views for the various sessions (Fig 4.A). Filters change the opacity of the brain surface to allow viewing of the cortical thickness of the brain and the lesions inside the brain.

Multi-sample community comparison. This tabbed view helps analyze the electrical spike propagations, the electrode communities, and their resulting inter-regional patterns across the multiple recorded sessions (T7) (Fig 2). The community helps identify those electrodes that spike together. Sessions are displayed side by side. As the dDTF algorithm generates many propagations between electrodes, neuroscientists note that it is almost impossible to gain insight by viewing the whole network. In accordance, we filter the network based on different percentile values of electrode propagations across the electrode activations. A higher percentile signifies high electrical activity.

Electrodes are colored-mapped by community groups. Electrodes not present, spiked, or below the percentile threshold are white. The network edges are shown as lines within a patch, and as raised arcs inter-patches, to reduce crossings. They are color-mapped to a gradient, where dark/light shade indicates the connection source/target. The width of the edge is mapped to the number of connections for that pair of electrodes, although thickness variations become negligible when filtered. Since overall clutter is massively reduced through the social network abstraction, the remaining density (clutter) in the network view serves as a client-desired cue to highly activated communities, and does not hinder analysis.

Multi-session electrode activation pattern conservation. This view helps the analysis of activated electrode regions across multiple sessions, indicating conserved patterns (Sec. 3.4.4) (T7). Areas activated in each session are marked in light grey, while conserved areas across sessions are marked in pink (Fig 4.B).

Multi-session patch comparison. This tabbed view helps analyze the brain network patterns across multiple sessions (T1). Furthermore, collaborators wished to identify similar network activity across brain regions (Fig 5.B) and patches (Fig 4.A), as well as the seizure onset (T4, T5). In addition, the view can help clinicians identify patterns of abnormal brain activity associated with specific brain lesions or seizures (T6). We show the network electrical activity of the electrodes, given a user-specified percentile threshold. Electrodes are color-mapped according to their patch/region group, and the edges use the same encodings as in the earlier views.

Electrode propagation over time. This tabbed view helps identify the source (onset) and spread of the electrode activations over time (T2). Despite its documented shortcomings, the surgeons wished to see an electrode activation animation to identify whether a seizure's cause is related to lesions for a selected time interval. Activated electrodes are highlighted through color. The size of the electrodes is mapped to the number of times an electrode spiked within that

time range. In addition, the propagation animation can be played in one-minute increments (also a surgeon request) across different percentile thresholds (see Supplemental Materials).

4.2. Brain Patch Network View

This view aids in exploring within-patch (Fig 4.C) or within-region (Fig 5.B) electrode networks (T4). Our collaborators wished for this view because high electrical activity patches indicate possible seizure onset. A 2D node-link diagram encodes this activity via the 2D snapshot extraction algorithm, which maintains the electrode spatial order for better visibility. Electrodes can be color-mapped according to their patches, regions, or community groups. The diagram nodes and edges use the same encodings and thresholding as the 3D view, and the views are linked through brushing and linking.

4.3. EEG Panel View

This view supports the exploration of EEG frequencies over time (T3) (Fig 4.D), which can be useful when determining which electrodes may contain seizure information. The EEG data contain frequencies at the millisecond level, so the file size is huge (each patient session is $>$1GB with 600,000 data points in a 10-minute recording). To solve the delays and errors due to this problem, we fetch the data on the back-end based on the user-selected patient sample and time range. However, this did not solve the loading time. To this end, we first converted our Comma Separated Values (CSV) file containing the EEG data to a fast parquet file, a standard binary data store for big data. This reduced the loading time from $\sim$12min to $\sim$5min. To further reduce this loading time to seconds, we incorporated a database management system, duckDB [RM19], to store and fetch the EEG data. This resolved the loading time issue, and data could be fetched and rendered to the interface within seconds. This view is also linked to the other views.

4.4. Patch Summary View

This view supports the analysis of localized patterns of spike frequencies, onset, and spread (T2) (Fig 4.E) with three similar but functionally different electrode distribution diagrams. In the spike onset summary diagram, we show the number of outgoing edges per electrode within the 10 minutes of the current sample by using basic color and size mapping of circular markers, in agreement with an existing visualization published in the epilepsy literature (Fig 6). In the spike spread summary diagram, we show in the same manner the number of incoming edges per electrode (Fig 6). In the patch/region summary diagram (Fig 4.E), we build towards a composite encoding based on Kosara's circular part-to-whole encodings, which have been found to be as effective in terms of error as linear bar charts [Kos19]. This encoding enables us to preserve the circular, compact layout preferred by our collaborators. For each electrode, a circle is divided into colored arc areas, then banked at 45% for improved perception. The first area is mapped to the number of spikes in this electrode, which did not contribute to propagation activities. The second area is mapped to the number of onset spikes at this electrode that generated outgoing propagation edges. The third is mapped to the number of spread spikes that generated incoming propagation edges at this electrode. The size of the circle is mapped to frequency, the total number of spike events at that electrode. Using these three diagrams, this view allows observing the relationships among spike frequency, participation in propagation, and signal onset/spread for each patch.

Our computational back-end was built using Python with Flask, pandas, scikit-learn, NumPy, and Jupyter. The front end was developed using JavaScript with D3.js [BOH11], three.js, and React libraries.

5. Evaluation and Results

In this section, we present both algorithmic validation and qualitative evaluation, including case studies and expert feedback from multiple clinical sites. We first evaluate our community algorithm on a synthetic dataset with known ground truth. We then evaluate the system via multiple demonstrations involving two neuroscience researchers who are also co-authors of the paper. One of these researchers (BM) was our main contact point for the project and consistently provided model development support and qualitative feedback. Together, we performed two case studies remotely using the think-aloud method, screen sharing, and note-taking. BM guided the exploration of the interface while the first author piloted. Furthermore, we conducted a task-oriented usability study with seven participants to assess system usability and effectiveness (see supplemental materials)

In addition to the regular design feedback sessions, we demonstrated the functionality of our solution to a group of 10 neurologists and surgeons at four medical centers: two neurologists and two epilepsy neuroscientists (University of Illinois Chicago), one neurology clinician specializing in epilepsy and one neuroscientist (Rush University), one neurologist specializing in the medical and surgical treatment of epilepsy and one neuroscientist (University of Chicago), one clinical coordinator for epilepsy surgery and one researcher (Northwestern University), and one pediatric neurology and epilepsy clinician (Northwestern's Lurie Children's hospital). The group members had extensive experience (10-29 years) in treating epileptic patients.

5.1. Algorithm Validation With Synthetic Data

In order to assess the efficiency of the modified dynamic community detection algorithm, we applied the algorithm to a synthetic dataset. The synthetic data is a dynamic network graph with communities that evolve over ten time steps in terms of both the number of communities and community membership. The initial network has three communities: com1, com2, and com3, with 9, 9, and 15 members, respectively. At time step three, two members are added to com1. At time step four, two members are removed from com3. At time step five, an additional community, com4, consisting of 10 members, is introduced. At time step six, two members are added to com2. At time step eight, two members are moved from com3 to com4. Members within each community exhibit dense connections, with sparse connections outside their respective communities.

Nodes are positioned to replicate different brain regions. Node connections were randomly created, and we varied the density values from 0.05 to 0.4, reflecting different network connectivities,

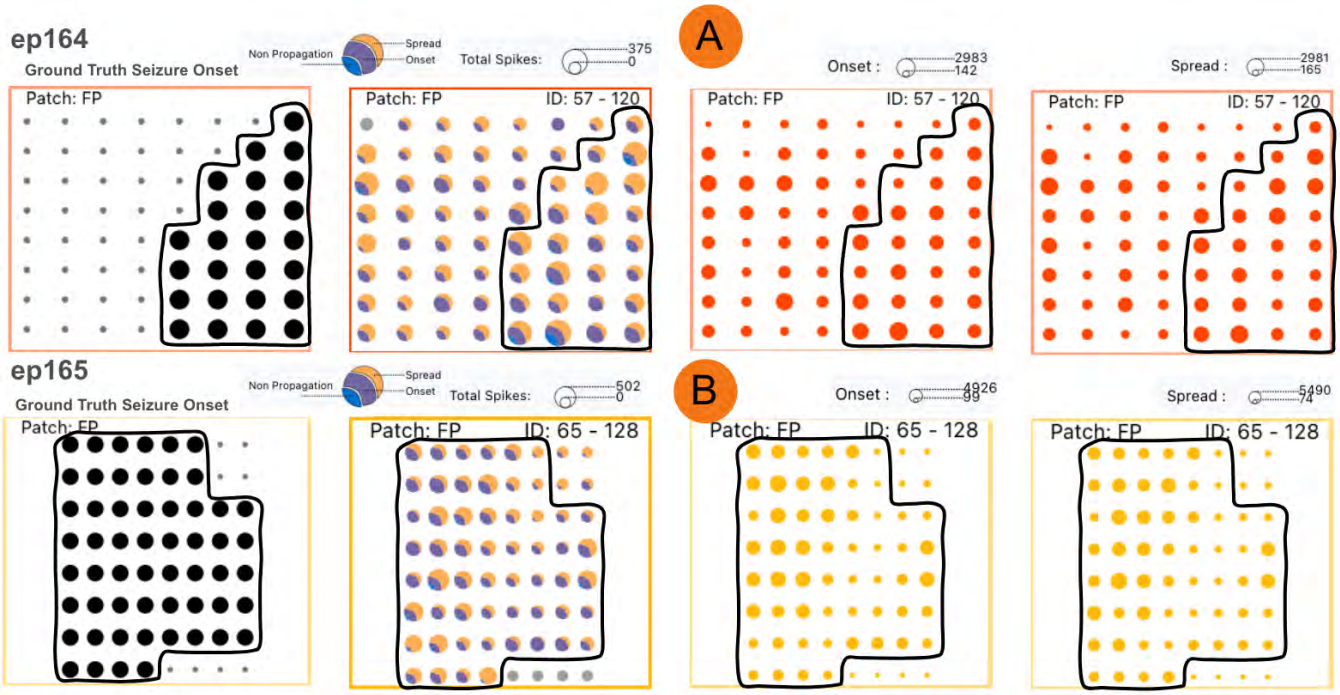

Figure 6: Electrode activation pattern exploration for patients ep164 and ep165 showing seizure onset, spike frequency, onset, and spread zones in the frontal-parietal (FP) region. The outlined region indicates the clinically identified seizure onset zone. Seizure onset zones are highlighted across all views. (A) ep164 correlates highly between seizure onset zones and spike frequency, onset, and spread zones. (B) ep165 does not have a high correlation. Many seizure electrodes have low spike frequency, onset, and spread.

then assessed the performance of the algorithm for each density value. We used this range as a density value greater than 0.4 yields dense connections across all members in the network, unlike in the application domain. We then applied the modified community detection technique for each density value and assessed accuracy, precision, recall, and f1 scores. The algorithm achieved 100% accuracy, i.e., correctly identified each community member across all time steps for density values between 0.05 and 0.3. The accuracy decreased to 88.57% accuracy for a high density of 0.4, and surprisingly, the lowest accuracy was 68.57% at a density value of 0.35. In these two edge cases, the algorithm failed to detect correct communities at time step 3 when two members were added to com1. However, we obtained high precision, recall, and f1 scores even in these two border cases, achieving a precision of 0.84, recall of 0.89, and f1-score of 0.85, even at a density of 0.35. Additionally, we achieved a precision of 0.995, recall of 0.996, and f1-score of 0.995 at density 0.4. These results indicate that the algorithm efficiently identifies community members, minimizing both false positives and false negatives. In addition, we also compared the results to those obtained using the original Louvain algorithm, which produced accuracies ranging from 42.85% to 100%. The original algorithm performed well at lower densities but poorly at higher densities, with an accuracy of 42.85% at a density of 0.4 and with a precision of 0.63, recall of 0.70, and f1-score of 0.66.

Finally, we tested our proposed algorithm in networks with different time steps ranging from 10 to 10000. At lower density values (between 0.05 and 0.25), the algorithm correctly identified community members in each time step with an accuracy of 99% to 100%. At higher density values, if the network remained static in the later timesteps, the algorithm tended to merge smaller communities into larger ones based on the network connectivity, leading to a decrease in overall accuracy. This confirms that the accuracy of the algorithm is independent of the number of time steps, but rather depends on the density of the network. On average, the algorithm run completed in less than two seconds for networks with up to 1000 timesteps. In a network with 5000 time steps, the completion time was ∼21 seconds, and in a network with 10000 time steps, the completion time was still acceptable, ∼50 seconds.

5.2. Ground-Truth Validation and Extended Exploration

Two patients (ep164 and ep165) had the site of the seizure onset zone identified by physicians as ground truth (outlined areas in Fig. 6 leftmost images). Our first aim was to replicate earlier dDTF raw analysis results performed within-patch, which found that the ground truth seizure onset zone was roughly correlated with spike onset and spread. Our second aim was to extend this exploration to multiple-patch analysis, which had not been performed before.

The neuroscientist loaded the data for ep164, a pediatric patient. He recalled that the physicians had marked the frontal-parietal (FP) region as the primary seizure onset zone but also indicated secondary onset zones in other areas. For the within-patch analysis, he examined the 3D view in conjunction with the patch summary view

(T4) and observed that patch #2 corresponded to the FP region. He also confirmed that the 2D patch extraction was correct in preserving electrode position and order. He confirmed that the physician-identified zone correlated with the number of spikes, spike onset, and spike spread (Fig. 6A).

Moving beyond single-patch analysis, he noted, in the 3D view, the propagations between the FP and lateral occipital (LO) region, as well as spike occurrences in the posterior subtemporal (PT) region (T2, T5) (Fig 5.A). This was in agreement with the physicians' note about secondary onset zones. Exploring other sessions, he observed similar spike occurrences. When examining the community view (Fig 5.A), he noticed these highly propagated electrodes formed four communities, conserving propagation patterns across the three sessions.

He then moved to the second ground-truth dataset (ep165), a case of infantile spasms. He repeated the steps, confirming that the physician-identified zone was again correlated with the total number of spikes, spike onset, and spike spread (Fig. 6.B). Moving beyond within-patch analysis, he observed in the 3D view and 2D network view high spike activity in the frontal-parietal (FP) area and secondary electrical activity in the posterior subtemporal (PT) and anterior subtemporal (AT) regions (T1, T2). The patch summary view (T4) supported this observation. He recalled these were indeed seizure onset regions, as identified by the surgeon. Moving on to the 3D patch network view, he noticed high, localized propagations within the FP region and in other areas (Fig. 5.B). He examined the result of the community detection algorithm across the three sessions available (Fig. 2) and noted a small group of electrodes inside the FP region is responsible for the high electrical activity across all sessions. These highly propagated electrodes were grouped into three communities, conserving the spike pattern (T7).

He then refocused on the FP region. He noticed that although most of the electrodes in this region corresponded to seizure onsets, few electrodes had high spike occurrence, spread, and onset (Fig. 6.B). He noted this was consistent with his previous analysis and that the seizure onset region encompassed a large portion of the grid. He selected several highly activated electrodes with network activities from the network view to see them in the EEG panel (T3). He noticed more spike occurrences than other electrodes, as expected. He noticed an inverse correlation between the electrodes with high spike occurrences and electrodes with high spike onsets. He remarked that, in many cases, brain regions with the highest spike occurrences do not initiate spiking. Overall, he noted that spike occurrences were unique to each patient, they were highly consistent, and the pattern was conserved across multiple sessions within each patient. He further remarked that the characterization of spike patterns and patches and the 2D extraction techniques were accurate and yielded results consistent with the surgeons' assessment.

5.3. Propagation Conservation Across Sessions

This case study investigated, in more depth, whether network propagation patterns are conserved across multiple sessions recorded for the same patient. The specific questions asked included: In a specific session, are there regions of high electrical activity? Are the spike onset regions correlated with the location of lesions? Are there specific propagation patterns (onset and spread) present? If so, are the patterns conserved over multiple sessions? Are there patterns only within regions or among various areas? Are the patterns conserved when using different filter thresholds? What is the electrical activity of the involved electrodes? Is the activity conserved across multiple sessions?

The exploration started by selecting and loading the data for an anonymized patient (ep129) from the patient list. This patient had two recorded sessions. The expert selected the 98th percentile threshold for the propagation network and observed that the dynamic community detection algorithm identified 2 communities in the first session and 3 in the second. He noticed that the location of the electrodes in the communities was largely conserved across the two sessions (Fig 4.B) and noted the anatomical location of the communities (T7). He also observed that most spike onset and spread regions were in the same patch in both sessions, with two spikes propagating to another patch (T2) (Fig 4.A). The expert then varied the percentile threshold and observed the electrode communities again (T7). He remarked that similar electrode communities were present across different percentiles, indicating high network activities, and that the electrode activation pattern was conserved across sessions, with a small number of spikes spread to other patches.

Next, he moved on to the 2D patch view (T2) (Fig 4.C), while frequently cross-checking with the 3D view (T2). Patch #2, located in the sub-frontal (SF) region, had more electrical activity and spike onsets than other patches. Filtering for the high electrical activity electrodes in the 2D patch view (T1, T6), he noted those electrodes were located near the brain lesion in the 3D view (T2) (Fig 4.A, B). He then used the EEG panel (T3, T2) (Fig 4.D) to confirm that these electrodes had more spike occurrences than the others. He confirmed the association with brain lesions by repeating the analysis on a second patient (ep187) with multiple brain lesions (see Supplemental Materials).

Last, he wished to analyze a patient with no lesions. He selected the ground-truth pediatric patient with spasms (ep165). Initial observation revealed that this patient had a larger propagation network than others (T4, T5), with high electrical activities in patch #0 or the frontal-parietal (FP) region. Most electrode spike onsets originated in this region, and a few spikes spread to other areas. Further analysis in the 3D patch view (T2) (Fig 5.B) revealed that a small group of electrodes inside the patch were responsible for the high electrical activity across all three sessions. He recalled that this patient used to have multiple electrical spasms, where the patient would contract their muscles and sit tight. Exploring the community view (T7), the expert found that, despite the lack of lesions, these highly propagated electrodes are still grouped into three communities across all sessions (Fig 2). Overall, he noted that the frontal-parietal region was the most common in terms of high electrical activity (T1). Further investigation validated this finding: the dataset consisted mainly of frontal lobe epileptic cases. Clinically, the conservation of spike patterns supported the surgeons' decision to resect the onset region for these patients.

5.4. Qualitative Clinician Feedback

During the group session, our solution yielded enthusiastic feedback, in particular from surgeons. Due to the minimal availability of the surgeons at the multiple medical centers, the demonstration was compact, and group feedback was limited to the meeting. Overall, the group repeatedly and spontaneously stated the implementation was "amazing," "very impressive work," "this is beautiful," "wow," and "really cool," both during the meeting and in a series of spontaneous emails after the meeting. They particularly appreciated and emphasized the potential value of the system in surgical treatment decisions, in particular the decision to pursue surgical resection. They noted the automated approach to detecting pattern conservancy was instrumental in making the decision process less subjective. They also appreciated the ability to pinpoint the specific brain regions that may be responsible for seizure onset, as surgically resecting these locations results in better seizure control in patients. They also noted that, in general, planning surgery for these cases was highly time-consuming but worth the effort, and that the system implementation was supporting a much-desired and essential functionality.

The neurologists in the meeting also remarked on the value of scaffolding the interface based on the state-of-the-art workflow: "When we change things, a physician could completely freak out. So we want [...] a new interface but keep important components similar so they can find similarities in what they have been seeing." They noted surgeons would most likely use the 3D view and conservation analysis results, typically at the 98-99% threshold, in addition to the EEG panel. The remaining views serve the neurologists' analysis and help increase confidence in the conservancy analysis. BM noted, "Some engineering-oriented physicians will love it, and [thanks to the scaffolding] clinical-oriented physicians will be okay with it."

5.5. Discussion

2nd generation ensembles. This was the third large project in our group where ensemble analysis required significant data transformation and feature extraction before meaningful comparison was possible. The earlier two projects [L*18, N*23] involved computational fluid dynamics data, whereas the present work focuses on measured data. Notably, In all these projects, once the relevant structures were extracted, many of the traditional challenges associated with classical ensemble visualization such as visual clutter and rendering complexity, were greatly reduced.

Legitimate theoretical observations do arise from practical applications, and recording such observations helps enrich the data visualization field [MM20]. To understand whether this extraction-first approach extended beyond our own work, we explored ensemble-visualization research across major visualization venues. We observed multiple examples, spanning fluid mechanics, climate, and biomedical imaging, where the key features were not directly available in the raw data and instead had to be derived before summarization or comparison. These works, highlighted in our Related Work, similarly required significant data transformation and feature extraction to make ensemble members meaningfully comparable.

Like our project, these applications involve complex problems or models, where the relevant features are "hard to define, but recognizable when seen" [L*18]. Several of them were also generated through direct measurement rather than simulations [v*23, C*20, N*25], often collected across multiple sites and teams. This results in workflows that merge domain expertise, data transformation, and feature extraction before any visual or statistical aggregation proves useful.

Motivated by this recurring pattern, we use the term "second-generation ensembles" to describe ensemble analysis challenges in which meaningful features must be defined or extracted before effective visualization, statistical analysis, or machine-learning approaches can be appropriately applied. This emergence of 2nd generation of ensemble datasets presents new challenges compared to first-generation models, particularly regarding data attributes and client needs. These datasets require innovative data transformation and feature extraction techniques in visual computing, which help mitigate rendering challenges faced by earlier ensembles. The process of identifying key features is often interactive and requires input from domain experts. As multi-site and multi-disciplinary studies become more common in life sciences, adapting traditional methods to assess these new datasets can result in inefficient design and evaluation cycles.

Solution merits and generalizability. The spatial methodology developed in this project, specifically the dense social network abstraction and transformation and subsequent spatial pattern extraction and conservancy, helped us create novel, desired functionality in the application domain. Our pattern conservation approach automates a largely subjective, manual, and cumbersome analysis process. Although each patient contributed 2–5 recording sessions (ranging from 120 to 600 seconds per session), this number of sessions is consistent with prior clinical studies [M*19] where it was shown to be sufficient for capturing stable within-patient activation patterns, as the spatial distribution and propagation of spikes tend to remain consistent for a given individual. Since seizure onset regions are specific to each patient, our objective was to validate the method's ability to accurately identify and extract activation regions rather than to generalize findings across multiple patients. Nonetheless, the proposed algorithm is scalable and can be applied to larger datasets and cohorts. Our dynamic community analysis and temporal Louvain approach may apply to other problems where spatial stability is important, such as drug design in biology [L*14] or finger formation in fluid mixes in CFD [L*18]. The dynamic community analysis algorithm can identify communities for any time-varying network.

Likewise, our patch 2D projection algorithm automates a manual, time-consuming layout process used in the application domain. The 2D snapshot extraction algorithm introduces an elegant way to project brain patches to 2D space, maintaining the sample spatial orientation, leading to improved visual clarity. This may be helpful in other spatial problems where visual clutter is present, beyond brain data, for example, dark matter formation in astronomy [HPAM19] or lymph node projection in head and neck oncology [L*20]. In terms of further generalizability, even though our visual computation solution aims to analyze domain-specific epileptic spike propagation, it can be used to explore other types of brain functional connectivity, for example, resting-state functional con-

nectivity. While the interface design aligns with established clinical workflows, these domain-specific requirements limit its direct adaptability to unfamiliar datasets. Extending the system to other domains would mainly require minor data remapping to fit the existing component structure, as the underlying analytical workflow for feature extraction, dynamic community detection, and pattern conservation remains unchanged. Finally, although the modified community detection and 2D projection algorithms were developed based on domain-specific data, they were implemented using standard and openly available Python libraries (e.g., NetworkX, scikit-learn, NumPy), with all parameters and steps thoroughly documented for reproducibility. These methods can therefore be reproduced by researchers with general computational expertise, without requiring deep specialization in visual computing and/or epilepsy research. Overall, this work presents a general strategy for analyzing complex, time-varying spatial data. By combining community detection with spatial abstraction, the approach identifies stable patterns within dense and dynamic networks. These ideas can guide visual computing efforts in other domains that face similar challenges, where understanding structure and change over time is key.

Scalability, assumptions and limitations. Our abstraction and feature extraction approach already reduces visual clutter, with remaining dense pockets providing a desired cue to high, concentrated electrical activity. Furthermore, our views are designed to handle data dynamically. The multi-tabbed comparative views can effectively display multiple brain networks without overwhelming visual clutter and accommodate more, depending on display size. Similarly, the brain patch network view and patch summary view can handle many brain regions, grouping many electrodes. To minimize edge crossing, we utilized a combination of lines and arcs. In addition, our back-end can handle large EEG data, while the front-end is able to access and render the data within seconds. The interface design intentionally balances complexity and familiarity by aligning visual layouts with existing clinical workflows. Most encodings follow established practices for EEG and neurosurgical visualization, while customized ones have been introduced through visual scaffolding to maintain interpretability. Threshold adjustments are also standard in clinical analysis. The multi-view interface integrates workflows that are typically accessed through separate tools (3D anatomy, EEG signals, and network connectivity), enabling users to interpret them together within a single workspace. Although multiple tabs and the patch summary view may require a brief learning period, the familiar visual structure helps minimize cognitive load. Animation is used selectively to display electrode activations over time, while key analytical results, such as community and pattern conservation analyses, remain static for long-term comparison. Finally, clinicians typically work with 10 to 20 seconds of EEG data and a limited set of electrodes, and our system includes filtering options to focus on relevant subsets, effectively reducing visual clutter and screen-space constraints.

The dynamic community detection algorithm can handle large networks. Depending on the network density and number of steps, it can also detect communities within seconds to a minute. Similarly, the electrode patch extraction algorithm is designed to automatically process large amounts of data. Additionally, as shown in Sec. 5.1, accuracy remains high across most density ranges, although there is a noticeable drop at a specific intermediate density of 0.35. This variation likely arises from the stochastic nature of community assignment when new members are added to highly connected networks, rather than indicating a systematic failure of the algorithm. Importantly, both precision and recall remained high, suggesting that the overall community structure was maintained even under these conditions. The electrode 2D snapshot extraction algorithm performs accurately on matrix or grid-based representations. Our method is learning-free and patient-specific, although recording conditions may vary. EEG recordings are inherently susceptible to noise and artifacts, which can impact spike detection and activation estimation. In this study, noisy or artifact electrodes were identified and removed by clinical collaborators prior to analysis, and the dDTF-based spike detection method was applied only to the cleaned signals (see Maharathi et al. [M*19] for detailed steps). Manual adjustments to electrode alignment and patch handling were limited to a few isolated instances and were visually verified to ensure accuracy, thereby minimizing potential bias. Although the cohort size and the number of sessions per patient are limited, these constraints are typical in clinical EEG studies and do not affect the validity of the proposed methodology.

There are also several assumptions and limitations to the current design of our solution, linked to the application domain data and tasks. Specifically, the visual encoding design, overall layout, number of sessions viewable on a regular screen, and solution implementation are built on specific domain knowledge and literacy, and domain data constraints. Per the surgeons' request, we also used an animation feature, which can hamper longer observations due to its reliance on short-term memory. Furthermore, the automated patch extraction relies on K-Means, which offers a simple and reproducible grouping but may not fully capture anatomically irregular structures. Likewise, although orthographic projection effectively preserves the overall spatial order for clinical interpretation, any 3D-to-2D mapping can introduce minor distortions. Future work includes implementing multi-patient electrical activity comparison, as well as conducting a more structured, task-based evaluation with multiple domain experts to assess usability and effectiveness of the system. Finally, due to human subject regulations, the existing system is currently available to our multi-site collaborators. As more data becomes available, we intend to integrate additional patient datasets and release the system for wider community use. We also plan to incorporate larger patient datasets to better quantify variability and evaluate reproducibility across different patients and recording conditions.

6. Conclusion

This work introduces a visual computing solution to automatically and objectively extract meaningful spatial measures of pattern conservancy across multiple ensemble members in dynamic epilepsy data. From a visual computing perspective, in addition to feature extraction, our solution introduces a social network-inspired dynamic network abstraction, a temporally-stable community detection algorithm, as well as registration algorithms to cross-correlate the spatial data, and to project the data to 2D for analysis. A front-end interface leverages custom encodings to support visual analysis of the data and model. Evaluation on synthetic and real-world data, and qualitative feedback from a multi-disciplinary team of experts

demonstrate the merits of this approach in aiding the investigation of epilepsy data before surgery. Last, we examined the analysis of spatial pattern conservancy in the frame of 2nd generation ensemble datasets, and discussed their defining characteristics and visualization challenges behind this emerging problem.

Acknowledgments

Our work is supported by NIH awards UG3-TR004501 and NCI-R01-CA258827, NSF award CNS-2320261, and the UIC Institute for Health Data Science. We thank our collaborators at four clinical sites, as well as the Electronics Visualization Laboratory, in particular Fabio Miranda for pointers to other 2nd gen ensemble works.

Ethics, Conflict of Interest, and Data Availability Statement

We declare that there are no conflicts of interest with this work, and this work adheres to established standards of research ethics and integrity. The data are not publicly available due to ethical and privacy concerns related to sensitive clinical information. The de-identified patient data shown in this manuscript were collected with consent under the UIC Institutional Review Board approved study STUDY2015-0457.

References

[A*08] AIGNER W., ET AL.: Visual methods for analyzing time-oriented data. *IEEE Trans. Vis. Comp. Graph. 14*, 1 (2008). 2

[A*13] ALPER B., ET AL.: Weighted graph comparison techniques for brain connectivity analysis. In *Proc. SIGCHI Conf. Hum. Factors Comp. Syst.* (2013), Assoc. Comp. Machinery. 2

[AA06] ANDRIENKO N., ANDRIENKO G.: *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Springer, 2006. 2

[AA11] ADRIENKO N., ADRIENKO G.: Spatial generalization and aggregation of massive movement data. *IEEE Trans. Vis. Comp. Graph. 17*, 2 (2011). 2

[AE*18] ABD EL—SAMIE F., ET AL.: A review of EEG and MEG epileptic spike detection algorithms. *IEEE Access 6* (2018). 1

[AFM06] AKIBA H., FOUT N., MA K.: Simultaneous classification of time-varying volume data based on the time histogram. In *Proc. Eighth Joint EuroGraph. IEEE VGTC Conf. Vis.* (2006), EuroGraph. Assoc. 2

[AG10] AYNAUD T., GUILLAUME J.: Static community detection algorithms for evolving networks. In *Int. Symp. Modeling Optimization Mobile, Ad Hoc, Wireless Networks* (2010). 4

[B*15] BOCK A., ET AL.: Visual verification of space weather ensemble simulations. In *IEEE Sci. Vis. Conf. (SciVis)* (2015). 2

[BGLL08] BLONDEL V., GUILLAUME J., LAMBIOTTE R., LEFEBVRE E.: Fast unfolding of communities in large networks. *J. Stat. Mechanics: Theory and Experiment 2008*, 10 (2008). 4

[BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D^3 data-driven documents. *IEEE Trans. Vis. Comp. Graph. 17*, 12 (2011). 9

[BPF14] BACH B., PIETRIGA E., FEKETE J.: Visualizing dynamic networks with matrix cubes. In *Proc. SIGCHI Conf. Hum. Factors Comp. Syst.* (2014), Association for Computing Machinery. 2

[C*20] CHEN J., ET AL.: Measuring the effects of scalar and spherical colormaps on ensembles of dmri tubes. *IEEE Trans. Vis. Comp. Graph. 26*, 9 (2020). 2, 12

[CP78] CARLBOM I., PACIOREK J.: Planar geometric projections and viewing transformations. *ACM Comput. Surv. 10*, 4 (1978). 6

[CS99] CAI W., SAKAS G.: Data intermixing and multi-volume rendering. *Comp. Graph. Forum 18*, 3 (1999). 2

[D*11] DOW S., ET AL.: Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Trans. Comput.-Hum. Interact. 17*, 4 (2011). 7

[dA01] DE CURTIS M., AVANZINI G.: Interictal spikes in focal epileptogenesis. *Prog. Neurobio. 63*, 5 (2001). 1, 3

[DMK05] DYKES J., MACEACHREN A. M., KRAAK M.-J. (Eds.): *Exploring Geovisualization*. Elsevier, 2005. 2

[DPJP20] DAHSHAN M., POLYS N. F., JAYNE R. S., POLLYEA R. M.: Making Sense of Scientific Simulation Ensembles With Semantic Interaction. *Comp. Graph. Forum 39*, 6 (2020). 2

[E*12] ENGEL J., ET AL.: Early surgical therapy for drug-resistant temporal lobe epilepsy: A randomized trial. *JAMA 307*, 9 (2012). 1

[F*17] FUJIWARA T., ET AL.: A visual analytics system for brain functional connectivity comparison across individuals, groups, and time points. In *IEEE Pac. Vis. Symp. (PacificVis)* (2017). 2

[F*19] FAVELIER G., ET AL.: Persistence atlas for critical point variability in ensembles. *IEEE Trans. Vis. Comp. Graph. 25*, 1 (2019). 1, 2

[FBW16] FERSTL F., BÜRGER K., WESTERMANN R.: Streamline variability plots for characterizing the uncertainty in vector field ensembles. *IEEE Trans. Vis. Comp. Graph. 22*, 1 (2016). 2

[FKRW16] FERSTL F., KANZLER M., RAUTENHAUS M., WESTERMANN R.: Visual analysis of spatial variability and global correlations in ensembles of iso—contours. *Comp. Graph. Forum 35*, 3 (2016). 2

[G*13] GRAMFORT A., ET AL.: MEG and EEG data analysis with mne—python. *Frontiers in Neurosci. 7*, 267 (2013). 6

[GN02] GIRVAN M., NEWMAN M. E. J.: Community structure in social and biological networks. *Proc. Nat. Academy Sci. 99*, 12 (2002). 4

[GZYL23] GUO Z., ZHAO Z., YAO L., LONG Z.: Improved brain community structure detection by two-step weighted modularity maximization. *PLOS ONE 18*, 12 (2023). 2

[H*11] HE B., ET AL.: eConnectome: A MATLAB toolbox for mapping and imaging of brain functional connectivity. *J. Neurosci. Methods 195*, 2 (2011). 2

[H*13] HÖLLT T., ET AL.: Visual analysis of uncertainties in ocean forecasts for planning and operation of off-shore structures. In *IEEE Pac. Vis. Symp. (PacificVis)* (2013). 2

[H*14a] HARENBERG S., ET AL.: Community detection in large-scale networks: a survey and empirical evaluation. *WIREs Comp. Stat. 6*, 6 (2014). 4

[H*14b] HÖLLT T., ET AL.: Ovis: A Framework for Visual Analysis of Ocean Forecast Ensembles. *IEEE Trans. Vis. Comp. Graph. 20*, 8 (2014). 2

[H*15] HASSAN M., ET AL.: EEGNET: An open source tool for analyzing and visualizing M/EEG connectome. *PLOS ONE 10*, 9 (2015). 2

[HE99] HEALEY C., ENNS J.: Large datasets at a glance: combining textures and colors in scientific visualization. *IEEE Trans. Vis. Comp. Graph. 5*, 2 (1999). 2

[HHB16] HAO L., HEALEY C., BASS S.: Effective visualization of temporal ensembles. *IEEE Trans. Vis. Comp. Graph. 22*, 1 (2016). 2

[HHK12] HOSSEINI S., HOEFT F., KESLER S.: GAT: A graph-theoretical analysis toolbox for analyzing between-group differences in large-scale structural and functional brain networks. *PLOS ONE 7*, 7 (2012). 2

[HPAM19] HANULA P., PIEKUTOWSKI K., AGUILERA J., MARAI G. E.: Darksky halos: use-based exploration of dark matter formation data in a hybrid immersive virtual environment. *Frontiers in Robotics and AI 6* (2019). 12

[HS04] HOCHHEISER H., SHNEIDERMAN B.: Dynamic query tools for time series data sets: Timebox widgets for interactive exploration. *Info. Vis. 3*, 1 (2004). 2

[HSSK16] HERMANN M., SCHUNKE A., SCHULTZ T., KLEIN R.: Accurate interactive visualization of large deformations and variability in biomedical image ensembles. *IEEE Trans. Vis. Comp. Graph. 22*, 1 (2016). 2

[HVHM21] HUANG H., VALENCIA G., HERMES D., MILLER K.: A canonical visualization tool for SEEG electrodes. In *Int. Conf. IEEE Eng. Med. Bio. Soc. (EMBC)* (2021). 6

[JC15] JOBST B. C., CASCINO G. D.: Resective epilepsy surgery for drug-resistant focal epilepsy: A review. *JAMA 313*, 3 (2015). 1

[K*08] KEHRER J., ET AL.: Hypothesis generation in climate research with interactive visual data exploration. *IEEE Trans. Vis. Comp. Graph. 14*, 6 (2008). 2

[KBH04] KOSARA R., BENDIX F., HAUSER H.: Time histograms for large, time-dependent data. In *Proc. Sixth Joint Eurographics IEEE TVCG Conf. Vis.* (2004), Eurographics Association. 2

[KBL22] KAPPE C., BÖTTINGER M., LEITTE H.: Topology-based feature analysis of scalar field ensembles: An application to climate (change) analysis. *Comp. & Graph. 104* (2022). 1, 2

[KH13] KEHRER J., HAUSER H.: Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Trans. Vis. Comp. Graph. 19*, 3 (2013). 2

[KML99] KIRBY R., MARMANIS H., LAIDLAW D.: Visualizing multivalued data from 2d incompressible flows using concepts from painting. In *Proc. Vis.* (1999). 2

[Kos19] KOSARA R.: The Impact of Distribution and Chart Type on Part-to-Whole Comparisons. In *EuroVis* (2019), The Eurographics Assoc. 9

[KSCW18] KABIR E., SIULY, CAO J., WANG H.: A Computer Aided Analysis Scheme for Detecting Epileptic Seizure from EEG Data. *Int. J. Comp. Intel. Sys. 11*, 1 (2018). 3

[L*14] LUCIANI T., ET AL.: FixingTIM: interactive exploration of sequence and structural data to identify functional mutations in protein families. *BMC Proc 8 (Suppl 2) S3* (2014). 12

[L*18] LUCIANI T., ET AL.: Details-first, show context, overview last: supporting exploration of viscous fingers in large-scale ensemble simulations. *IEEE Trans. Vis. Comp. Graph. 25* (2018). 1, 2, 12

[L*20] LUCIANI T., ET AL.: A spatial neighborhood methodology for computing and analyzing lymph node carcinoma similarity in precision medicine. *J. of Biomed. Info. 112* (2020). 12

[M*08] MUIGG P., ET AL.: A four-level focus+context approach to interactive visual analysis of temporal features in large scientific data. *Comp. Graph. Forum 27*, 3 (2008). 2

[M*15] MA C., ET AL.: Visualizing dynamic brain networks using an animated dual-representation. In *Eurograph. Conf. Vis.* (2015). 2

[M*17] MA C., ET AL.: RemBrain: Exploring dynamic biospatial networks with mosaic matrices and mirror glyphs. *J. Imaging Sci. Tech. 61*, 6 (2017). 2

[M*19] MAHARATHI B., ET AL.: Interictal spike connectivity in human epileptic neocortex. *Clin. Neurophysio. 130*, 2 (2019). 1, 3, 13

[Mar15] MARAI G.: Visual Scaffolding in Integrated Spatial and Non-spatial Analysis. In *EuroVis Workshop Vis. Analyt. (EuroVA)* (2015). 4

[Mar18] MARAI G. E.: Activity-Centered Domain Characterization for Problem-Driven Scientific Visualization. *IEEE Trans. Vis. Comp. Graph. 24*, 1 (2018). 4

[MM20] MARAI G. E., MÖLLER T.: The fabric of visualization. In *Foundations of Data Visualization*. Springer, 2020, pp. 5–14. 12

[MWK14] MIRZARGAR M., WHITAKER R., KIRBY R.: Curve boxplot: Generalization of boxplot for ensembles of curves. *IEEE Trans. Vis. Comp. Graph. 20*, 12 (2014). 2

[N*23] NIPU N., ET AL.: Visual analysis and detection of contrails in aircraft engine simulations. *IEEE Trans. Vis. Comp. Graph. 29*, 1 (2023). 1, 2, 12

[N*25] NIPU N., ET AL.: Lessons from the development and deployment of an interactive oncological risk estimator. In *IEEE Workshop on Visual Analytics in Healthcare (VAHC)* (2025), pp. 29–35. 1, 12

[NBB17] NICOLINI C., BORDIER C., BIFONE A.: Community detection in weighted brain connectivity networks beyond the resolution limit. *NeuroImage 146* (2017). 2

[New06] NEWMAN M. E. J.: Modularity and community structure in networks. *Proc. Nat. Academy Sci. 103*, 23 (2006). 4

[OSR12] OLUIGBO C., SALMA A., REZAI A.: Deep Brain Stimulation for Neurological Disorders. *IEEE Reviews in Biomed. Eng. 5* (2012). 3

[P*09] POTTER K., ET AL.: Ensemble-vis: A framework for the statistical visualization of ensemble data. In *IEEE Int. Conf. Data Mining Workshops* (2009). 2

[PP14] PACHORI R., PATIDAR S.: Epileptic seizure classification in EEG signals using second-order difference plot of intrinsic mode functions. *Comp. Methods Prog. Biomed. 113*, 2 (2014). 3

[PWH11] PÖTHKOW K., WEBER B., HEGE H.: Probabilistic marching cubes. *Comp. Graph. Forum 30*, 3 (2011). 2

[R*15] RICHIARDI J., ET AL.: Correlated gene expression supports synchronous activity in brain networks. *Sci. 348*, 6240 (2015). 2

[RBG07] RAUTEK P., BRUCKNER S., GROLLER E.: Semantic layers for illustrative volume rendering. *IEEE Trans. Vis. Comp. Graph. 13*, 6 (2007). 2

[RC18] ROSSETTI G., CAZABET R.: Community discovery in dynamic networks: A survey. *ACM Comput. Surv. 51*, 2 (2018). 4

[RM19] RAASVELDT M., MÜHLEISEN H.: DuckDB: an embeddable analytical database. In *Proc. Int. Conf. Management Data* (2019), Assoc. Comp. Machinery. 9

[Rou87] ROUSSEEUW P.: Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comp. Applied Mathematics 20* (1987). 6

[S*05] SALVADOR R., ET AL.: Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral Cortex 15*, 9 (2005). 2

[S*10] SANYAL J., ET AL.: Noodles: A tool for visualization of numerical weather model ensemble uncertainty. *IEEE Trans. Vis. Comp. Graph. 16*, 6 (2010). 2

[S*22] SHI L., ET AL.: MVNet: Multi-variate multi-view brain network comparison over uncertain data. *IEEE Trans. Vis. Comp. Graph. 28*, 12 (2022). 2

[SA*10] SANZ-ARIGITA E., ET AL.: Loss of 'small-world' networks in alzheimer's disease: Graph analysis of fMRI resting—state functional connectivity. *PLOS ONE 5*, 11 (2010). 2

[SSWZ23] SUPRIYA S., SIULY S., WANG H., ZHANG Y.: Epilepsy detection from EEG using complex network techniques: A review. *IEEE Reviews Biomed. Eng. 16* (2023). 2

[T*21] TIAN X., ET AL.: An improved louvain algorithm for community detection. *Mathematical Prob. Eng. 2021* (2021). 5

[TSWS05] TOMINSKI C., SCHULZE-WOLLGAST P., SCHUMANN H.: 3D information visualization for time dependent data on maps. In *Int. Conf. Info. Vis.* (2005). 2

[v*23] VAN DIJK L., ET AL.: Head and neck cancer predictive risk estimator to determine control and therapeutic outcomes of radiotherapy (hnc-predictor): development, international multi-institutional validation, and web implementation of clinic-ready model-based risk stratification for head and neck cancer. *European Journal of Cancer 178* (2023). 1, 12

[W*19] WENTZEL A., ET AL.: Cohort-based t-ssim visual computing for radiation therapy prediction and exploration. *IEEE Trans. Vis. Comp. Graph. 26*, 1 (2019). 2

[W*24] WENTZEL A., ET AL.: Ditto: A visual digital twin for interventions and temporal treatment outcomes in head and neck cancer. *IEEE Trans. Vis. Comp. Graph. 31*, 1 (2024), 65–75. 2

[WCLE05] WORSLEY K., CHEN J., LERCH J., EVANS A.: Comparing functional connectivity via thresholding correlations and singular value decomposition. *Philos. Trans. R. Soc. B Biol. Sci. 360*, 1457 (2005). 2

[WHLS18] WANG J., HAZARIKA S., LI C., SHEN H.: Visualization and visual analysis of ensemble data: A survey. *IEEE Trans. on Vis. and Comp. Graph. 25*, 9 (2018). 1, 2

[WMK13] WHITAKER R., MIRZARGAR M., KIRBY R.: Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles. *IEEE Trans. Vis. Comp. Graph. 19*, 12 (2013). 2

[X*19a] XU K., ET AL.: Ensemblelens: Ensemble-based visual exploration of anomaly detection algorithms with multidimensional data. *IEEE Trans. Vis. Comp. Graph. 25*, 1 (2019). 2

[X*19b] XU R., ET AL.: TempoCave: Visualizing dynamic connectome datasets to support cognitive behavioral therapy. In *IEEE Vis. Conf.* (2019). 2

[XWH13] XIA M., WANG J., HE Y.: BrainNet viewer: a network visualization tool for human brain connectomics. *PloS one 8*, 7 (2013). 2

[Y*17] YANG X., ET AL.: Blockwise human brain network visual comparison using NodeTrix representation. *IEEE Trans. Vis. Comp. Graph. 23*, 1 (2017). 2

[YC23] YE Z., CHEN M.: Visualizing ensemble predictions of music mood. *IEEE Trans. Vis. Comp. Graph. 29*, 1 (2023). 2

[YLL10] YANG B., LIU D., LIU J.: Discovering communities from social networks: Methodologies and applications. *Handbook Soc. Net. Tech. App.* (2010). 4

[Z*17] ZHANG C., ET AL.: Overview + detail visualization for ensembles of diffusion tensors. *Comp. Graph. Forum 36*, 3 (2017). 2

[ZFB10] ZALESKY A., FORNITO A., BULLMORE E.: Network-based statistic: Identifying differences in brain networks. *NeuroImage 53*, 4 (2010). 2

Supplemental Materials
Neurosurgery Network Pattern Analysis with 2nd Generation Ensembles

N. Nipu et al.

1 Data Description and Experimental Results

Table 1: Propagation Network Description

Patient Id	**Sessions**	**# Electrodes (# nodes)**	**Network size (# edges)**
ep129	session 1	79	26916
	session 2	79	4264
ep187	session 1	114	11488
	session 2	114	34237
	session 3	114	48280
ep162	session 1	102	112792
	session 2	102	25924
	session 3	102	64646
ep164	session 1	120	139880
	session 2	120	184903
	session 3	120	207520
ep165	session 1	106	205932
	session 2	106	185285
	session 3	106	192715

Table 2: Experimental Results (Communities)

Patient ID	Session	Network size (per percentile)							# communities (per percentile)						
		99	98	97	96	95	90	50	99	98	97	96	95	90	50
ep129	session 1	1684	2808	3813	4703	5833	9103	22518	2	4	3	3	4	4	3
	session 2	206	281	380	550	550	892	2750	3	2	3	4	4	4	2
ep187	session 1	983	1584	1976	2551	2778	4111	8985	3	3	2	3	2	3	2
	session 2	2238	4084	5394	6926	7948	12936	29043	2	2	3	3	3	3	2
	session 3	3365	5728	8225	10053	11998	19878	19878	3	3	3	3	4	2	2
ep162	session 1	7949	13735	18706	23093	27309	44113	99348	3	3	3	3	3	3	4
	session 2	1645	2891	4059	5026	5812	9509	21964	3	3	3	3	2	2	3
	session 3	4600	7688	10490	12835	15026	23607	53878	3	2	3	3	3	2	2
ep164	session 1	6194	10649	14854	19325	22000	38964	110605	7	8	5	4	3	3	3
	session 2	9259	16194	21168	27508	32790	54938	150922	5	5	4	4	3	3	3
	session 3	9446	17034	22429	28267	33952	57550	165354	4	4	3	3	4	4	3
ep165	session 1	8829	15855	22320	28639	34864	60139	166076	5	3	5	5	3	3	2
	session 2	7797	13010	18412	23303	28886	49208	145037	4	3	4	3	4	2	2
	session 3	7112	13367	18750	24140	28504	50689	50689	4	4	4	4	4	2	2

2 Task-Oriented Usability Study

2.1 Study Design

Surgeons, including neurosurgeons, neuroscientists, and other time-pressured clinicians, have no availability for user studies due to their punishing schedule. However, we conducted a task-oriented usability study to evaluate how effectively users could perform clinically motivated analyses using our system. Seven participants (E1-E7) took part in the study. Each participant completed a set of six tasks that directly corresponded to the analytical features our system is designed to support. These tasks included identifying spike-dense regions, analyzing propagation patterns across inter- and intra-patches, correlating propagation with lesions, verifying EEG data, and assessing the conservation of multi-session patterns. The study followed a specific procedure: participants were first given a brief overview of the interface, after which they completed all six tasks in order and at their own pace. We recorded the time taken to complete each task. Upon completing all tasks, participants answered three difficulty-related questions for each task and completed a System Usability Scale (SUS) questionnaire, along with providing additional optional feedback.

2.2 Participants

A total of seven participants (E1-E7) were included in the study. The group consisted of seven researchers with varying degrees of experience working with biomedical and/or neuroscience data. While most participants were familiar with general visualization concepts, their backgrounds varied in relation to the specific domain knowledge of EEG or spike propagation analysis. One participant reported extensive prior experience working directly with EEG or spike-propagation data, while the remaining participants had little to no exposure to these areas. Participants also indicated whether they had previously worked with biological or biomedical datasets of any kind. Five participants had some prior exposure to biomedical or biological data, while two reported no such experience. This distribution allowed us to evaluate how effectively the system supports both users with domain familiarity and those

encountering these analytical tasks for the first time.

2.3 Tasks

Participants performed the following six analytical tasks:

- **T1. (Identify Spike Activity Regions and electrodes)** Use the 3D brain propagation view to identify the region(s) and electrodes with the highest concentration of spike events in Session 1 for patient ep165.
- **T2. (Determine High-Activity Patch)** Use the Patch Network View to identify the patch with the highest within-patch propagation.
- **T3. (Identify Cross-Patch Propagation)** Use the Multi-Session View to identify whether any cross-patch propagation exists above the 96th percentile threshold for patient Ep129 across sessions.
- **T4. (Relate Activity to Lesions)** For patient ep129, use the 3D Anatomical Overlay View to determine whether high-activity electrodes are located near the lesion region in the brain model.
- **T5. (EEG-Based Spike Verification)** Use the EEG Panel to select highly propagated electrodes and verify their spike pattern through the EEG recordings.
- **T6. (Assess Pattern Conservation)** For patient ep187, use the Conservation View to determine whether spike propagation patterns are conserved across sessions.

These tasks collectively required participants to navigate all major interactive views, interpret spatial, temporal, and network-based information, and derive insights across multiple linked views.

2.4 Task Completion Time and Difficulty Ratings

Across all participants, the mean completion time per task ranged from 23.77 to 54.73 seconds. Tasks that involved a single view, such as T2 and T4, were completed the fastest. In contrast, tasks requiring interaction with multiple views, like T5 and T6, naturally took more time. The average times for each task were as follows: T1 - 54.73 s, T2 - 23.77 s, T3 - 35.69 s, T4 - 29.22 s, T5 - 43.51s, and T6 - 50.81 s. Despite the analytical complexity, all tasks were completed successfully, and task completion times remained consistently under one minute on average. This indicates that users were able to learn and navigate the interface efficiently.

Difficulty ratings showed a similar trend. Tasks T1-T5 were generally rated as easy or moderately easy by most participants. T6, which required interpreting propagation patterns across multiple sessions, was also rated as easy by the majority of participants (scores of 4 or 5). However, one participant, who reported no prior EEG experience, rated T6 as difficult (score of 1), resulting in higher variability in the ratings for this task. This suggests that while the multi-session comparison view is largely accessible, it may benefit from additional onboarding support for users without prior domain familiarity. Notably, participants with no prior EEG or biomedical background performed similarly to those with experience, suggesting that the interface effectively supports users who are new to the field.

2.5 System Usability Scale (SUS)

We also collected feedback using the System Usability Scale (SUS) framework. Users were presented with twelve questions, each rated on a scale from 1 to 5, along with a final question assessing how likely they were to recommend the tool to their colleagues in the neuroscience field, rated on a scale from 1 to 10. Additionally, we included open-ended questions to obtain feedback on specific features and potential improvements.

The system achieved a SUS score of 90.0, characterizing its usability as "excellent." Quantitative feedback indicated that users felt confident using the system ($M = 4.71 \pm 0.49$) and praised its functional integration ($M = 4.71 \pm 0.49$). Participants reported that they would be likely to use the system when studying spike propagations ($M = 4.86 \pm 0.38$) and agreed that the interface was easy to use ($M = 4.14 \pm 1.07$), as well as that most users would be able to learn the tool quickly ($M = 4.57 \pm 0.53$). Conversely, participants generally disagreed that the system was overly complex ($M = 1.29 \pm 0.49$), inconsistent ($M = 1.14 \pm 0.38$), cumbersome ($M = 1.29 \pm 0.49$), or that substantial support would be needed to use it ($M = 1.86 \pm 0.90$). Additionally, participants expressed a positive overall impression of the system ($M = 4.57 \pm 0.53$) and indicated a high likelihood of recommending it to other neuroscience or neurosurgery researchers ($M = 9.29 \pm 0.76$ on a 10-point scale).

The system has received positive feedback for its interactive and visually clear design, which effectively explores spike propagation patterns. Participants praised the 3D multi-session propagation view for its interpretability and ability to generate insights quickly. However, some participants identified opportunities for refinement, particularly in the EEG panel, where they felt additional contextual support or instructions were needed. Overall, the evaluation indicates that the system demonstrates high usability and strong user acceptance, with the need for improvements primarily focused on domain knowledge rather than core interaction design.

3 Electrode Activation Pattern Conservation Summarization Algorithm

Procedure 1 Electrode Activation Pattern Extraction

function ExtractElectrodePattern(E, G, N, Q)
    E: Electrode list with 3D coordinates
    G: Anatomical 2D grids per region (row-column matrices)
    N: Electrode activation networks (per session)
    Q: List of quantile thresholds
    Output: V – 3D visual elements (meshes, lines, spheres)
    for all quantile $q \in Q$ **do**
        for all sample $s \in N$ **do**
            $N_q[s] \leftarrow$ Filter edges in $N[s]$ using quantile q
            $C \leftarrow$ Connected components from $N_q[s]$
            for all component $c \in C$ **do**
                $NOG \leftarrow \emptyset$ ▷ Set of non-overlapping groups
                for all region r **do**
                    $c_r \leftarrow$ electrodes in c within region r
                    $G_r \leftarrow$ Grid matrix for region r
                    $visited \leftarrow \emptyset$
                    for all electrode $e \in c_r$ **do**
                        if $e \notin visited$ **then**
                            Initialize $queue \leftarrow [e]$, $group \leftarrow \emptyset$
                            while $queue$ not empty **do**
                                Pop e_k from $queue$
                                Add e_k to $group$ and $visited$
                                for all 4-way neighbors of e_k in G_r **do**
                                    if neighbor $\in c_r$ and not in $visited$ **then**
                                        Add to $queue$
                                    end if
                                end for
                            end while
                            Add $group$ to NOG
                        end if
                    end for
                end for
                $main \leftarrow$ largest group in NOG
                for all group $g \in NOG$, $g \neq main$ **do**
                    Find nearest electrode pair (e_1, e_2) between g and others
                    $path \leftarrow$ line from e_1 to e_2
                    for all electrode $e \in E$ **do**
                        if e lies on $path$ **then**
                            Adjust path by inserting points with slight variation before & after e
                        end if

```
                end for
                Store adjusted path in V
            end for
            for all group g ∈ NOG do
                if |g| ≥ 4 then
                    Compute convex hull of 3D positions in g
                    Add hull mesh to V
                else if |g| = 1 then
                    Add sphere at 3D position to V
                else
                    Add line between electrodes in g to V
                end if
            end for
        end for
    end for
    return V
end function
```

4 Electrode 2D Snapshot Extraction

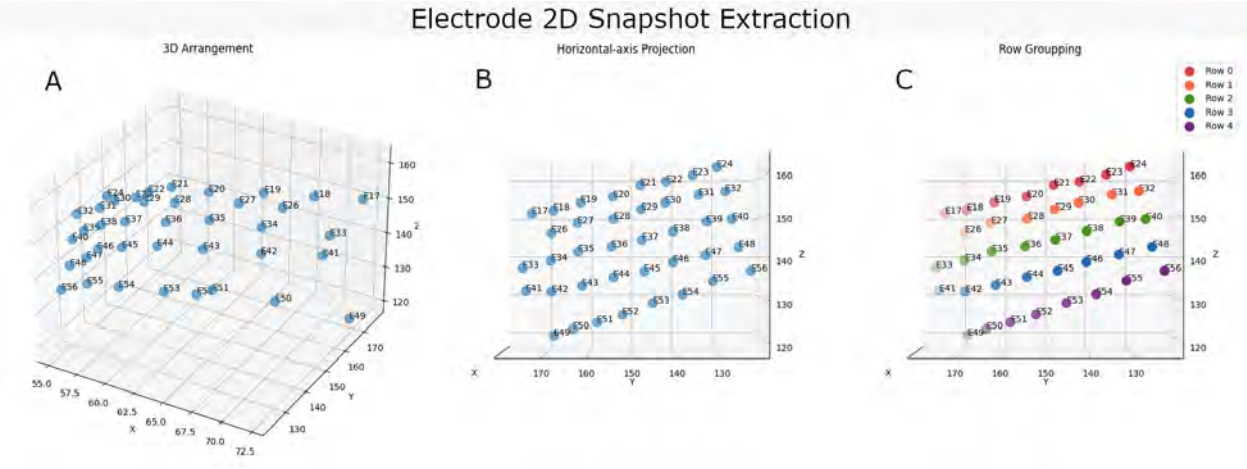

Figure 1: Electrode 2D Snapshot Algorithm: (A) Original arrangement of electrodes within a patch in 3D space (B) Result of the horizontal axis projection, arranging the 3D patch in the superficial (SF) region (C) Extracted 2D layout leveraging dynamic thresholding and Euclidean distances.

5 Electrode Patches

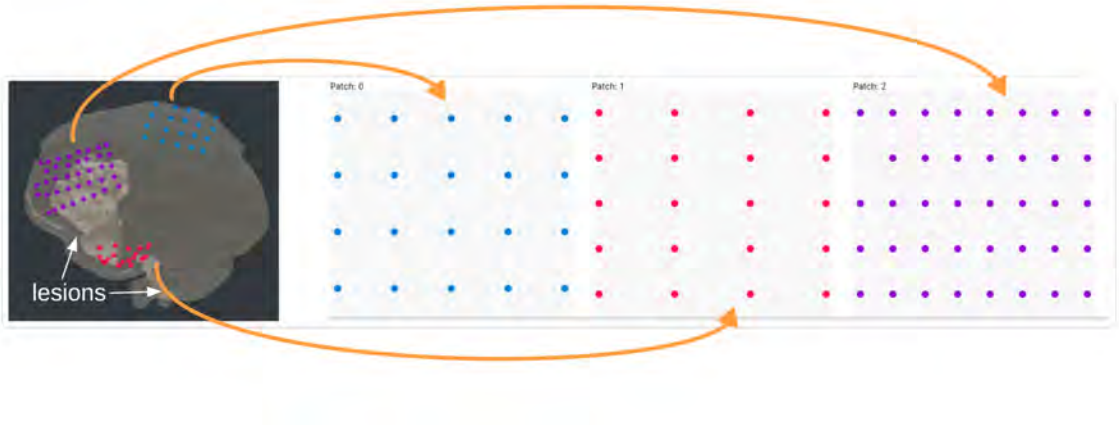

Figure 2: Electrode patches extraction and their 2D representations. Electrodes on the 3D brain surface are grouped into patches using the K-means-based patch extraction algorithm (left). The corresponding 2D layout snapshots (right) illustrate the spatial structure of each detected patch after projection, maintaining electrode orientation for within-patch and across-patch analysis.

6 Electrode Activated Region Exploration

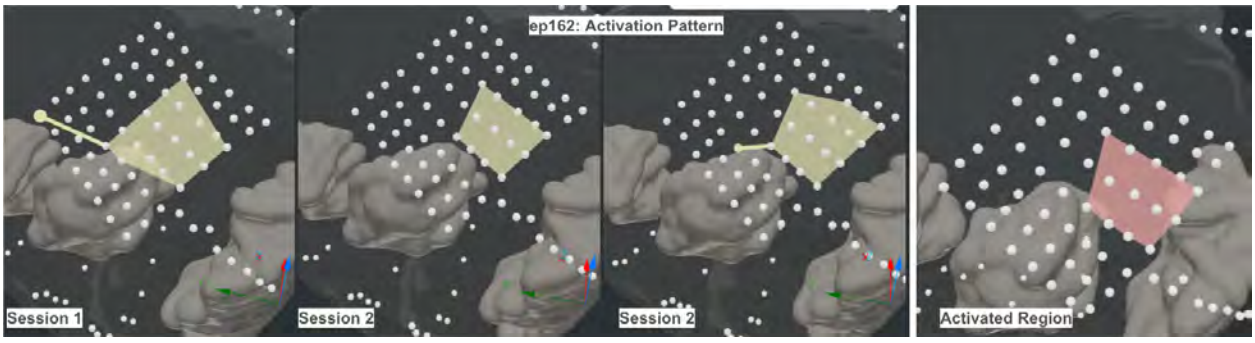

Figure 3: Electrode activated regions of patient ep 162, showing the activated regions for each session and the conserved activated region across all sessions.

7 Multi-region Spike Propagation and Activation Analysis

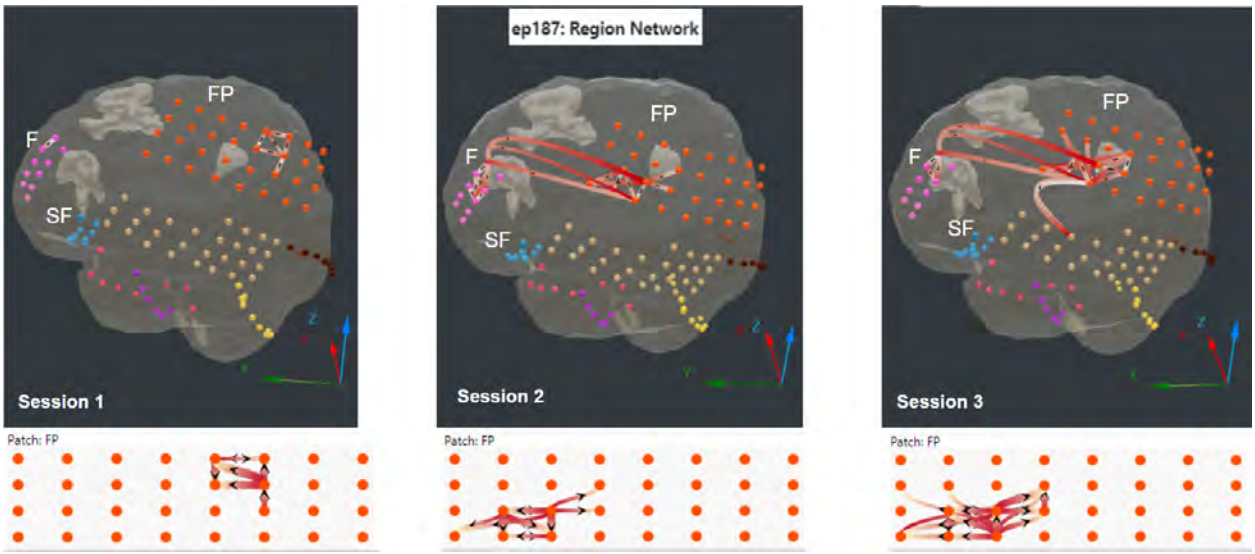

Figure 4: Multi-region spike propagation and activation analysis of ep187. Spike propagation across regions above the 99th percentile in three sessions, showing high electrical activities in FP and F regions near the three brain lesions and no electrical activity in the SF region (seizure onset zone). Within-patch network confirming high-electrical activity in patch FP

8 Electrode Propagation per-Minute Exploration

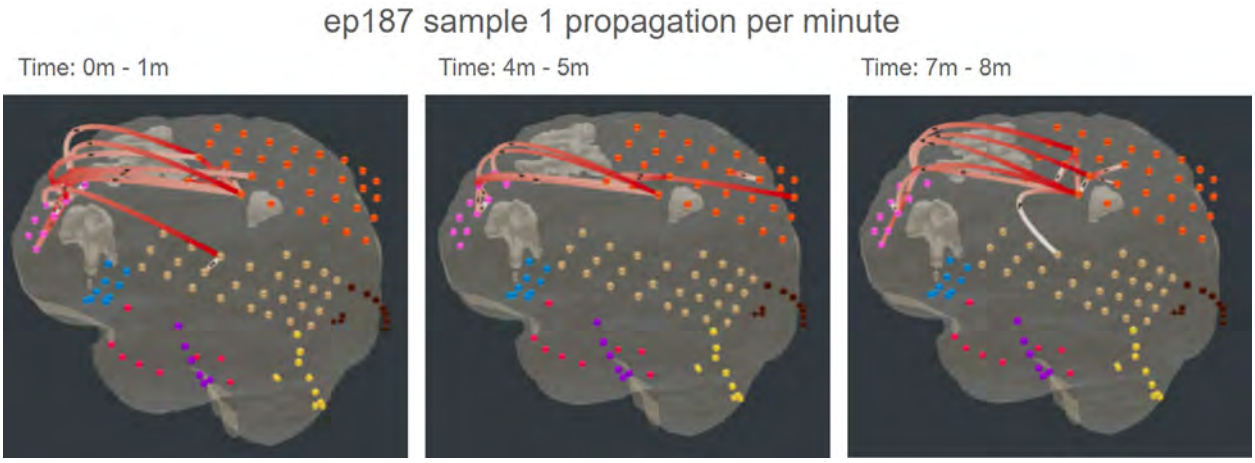

Figure 5: Electrode propagation per minute showing three time steps of ep187 sample 1. High network activities can be observed in frontal (F) and superior frontal parietal (FP) regions across all snapshots near the lesions, and some spikes spread between these patches.

9 Electrode Activation Pattern

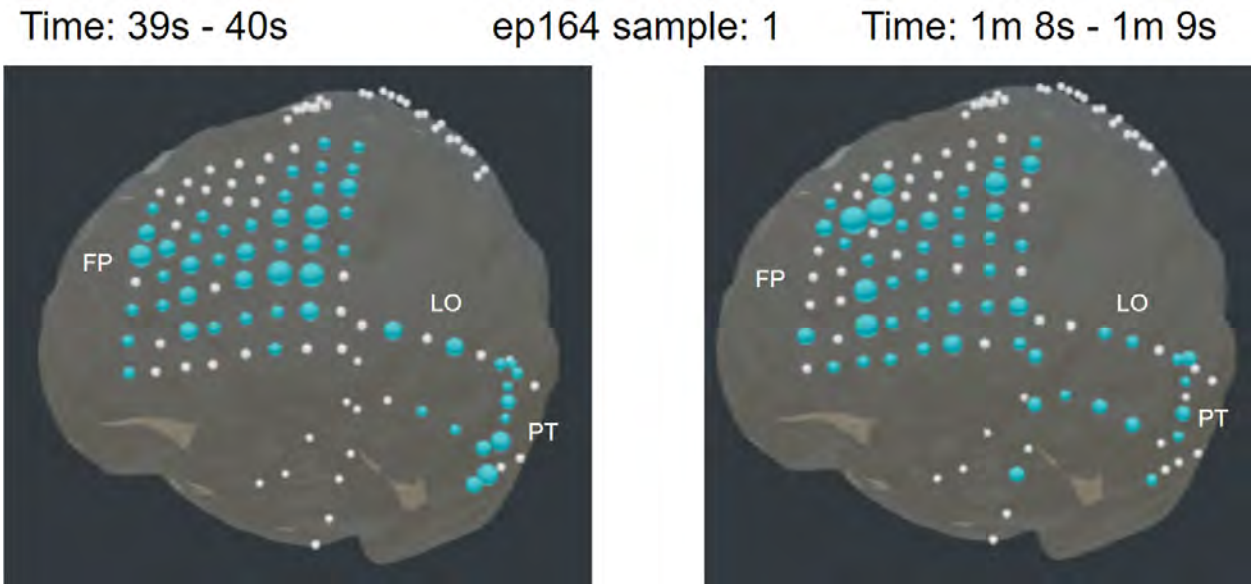

Figure 6: Electrode activation pattern exploration. Electrode propagation over time view showing activated electrodes per second in two-time steps in patient ep164 session one. The activated electrodes are colored where size corresponds to the number of spikes. In both time steps, high activation is observed in the FP and the PT regions.