soloncouncil : Pages:1–17

Original Research

Uncertainty-Aware Latent Neural-Field Modeling on Patient-Specific Graphs for Seizure Dynamics Estimation and Interventional Planning

Renato Alcaraz[1] and Marvin Guzman[2]

[1]University of the East, Aurora Boulevard, Quezon City 1108, Philippines.
[2]Jose Maria College Foundation, Ecoland Drive, Davao City 8000, Philippines.

Abstract

Stereo- and scalp-electroencephalography provide multichannel time series that partially observe complex, spatially distributed neural dynamics implicated in focal seizure generation and spread. Clinical decision-making increasingly seeks computational support that can handle sparse electrode coverage, nonstationary dynamics, and heterogeneous patient anatomy, while remaining robust to artifacts and modeling mismatch. This paper develops an uncertainty-aware framework that treats seizures as emergent transitions of a latent neural field evolving on a patient-specific graph, where nodes represent localized cortical or deep structures and edges encode directed, time-varying effective influence. We propose a stochastic latent neural-field model that couples continuous-time dynamics with discrete event modulation to capture pre-ictal drift, ictal recruitment, and post-ictal recovery within a single probabilistic program. Inference is performed via structured variational methods that yield calibrated posterior distributions over connectivity, latent states, and regime parameters, enabling principled uncertainty quantification for clinical interpretation. Building on these posteriors, we formulate intervention planning as a constrained optimal control problem over feasible ablation or stimulation operators, incorporating safety-aware penalties and sparsity constraints to encourage minimally disruptive strategies. The resulting algorithm produces individualized intervention scores, target sets, and confidence measures without assuming a single mechanistic cause of seizures. Experiments on retrospective intracranial and scalp recordings, complemented by anatomically informed simulations, demonstrate improved stability of inferred networks under subsampling and noise, as well as consistent ranking of intervention candidates across initialization and hyperparameter sweeps. The framework is designed to be modular, supporting different sensing modalities and operational constraints, and emphasizes transparent uncertainty reporting alongside predictive performance.

1. Introduction

Epileptic seizures are macroscopic phenomena arising from distributed neural interactions that unfold over multiple spatial and temporal scales [1]. The practical challenge in electrophysiological monitoring is that recordings provide only a partial, noisy projection of these interactions, with coverage shaped by clinical hypotheses, surgical considerations, and patient-specific anatomy. Scalp EEG offers noninvasive access but suffers from volume conduction, limited spatial specificity, and susceptibility to artifacts. Intracranial recordings, including depth electrodes, increase localization and signal-to-noise but remain spatially sparse relative to the underlying circuitry [2]. Computational models used in clinical contexts therefore face a tension between expressivity and identifiability: richer models can represent nonstationary propagation patterns, but their parameters may be weakly constrained by limited observations, increasing the risk of overfitting and unstable recommendations.

A second challenge is that the clinically relevant outputs are not limited to detecting seizures. For many workflows, the goal is to characterize where seizures originate, how they propagate, and which interventions are most likely to reduce seizure burden while preserving function [3]. This involves

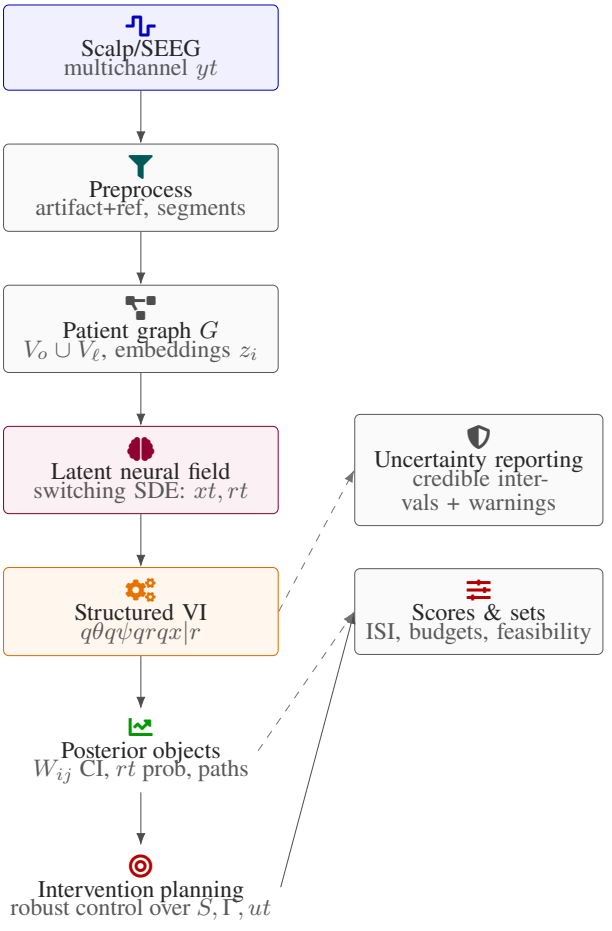

Figure 1: End-to-end pipeline: multichannel EEG/SEEG is mapped to a patient-specific directed graph and a switching latent neural-field SDE. Structured variational inference yields posterior distributions over regimes, trajectories, and effective influence, which are propagated into robust intervention planning to output target sets with confidence measures.

combining information across time windows, capturing directional interactions, and explicitly handling uncertainty due to short recordings, variable seizure morphologies, and incomplete sampling of the epileptogenic network. In addition, practical deployment requires that algorithms remain robust under patient-specific variability and operational constraints such as limited compute, missing channels, and differences in referencing schemes. Work on low-complexity monitoring has emphasized personalization and feature selection to tailor performance to individual subjects while controlling computational costs, exemplified by approaches that adapt channel and feature subsets per patient for wearable-friendly monitoring [4]. Such perspectives motivate modeling strategies that treat patient identity not as a nuisance variable but as a first-class component of the inference problem.

A third challenge concerns intervention planning [5]. Whether the intervention is resection, ablation, or stimulation, the choice of target regions is constrained by anatomy, safety, and functional preservation. A model that only ranks nodes by a single centrality metric may fail when seizures depend on distributed loops or when observed nodes are not the true drivers but rather downstream relays. Moreover, any ranking should be accompanied by confidence estimates that communicate when the data are insufficient to support a stable conclusion [6]. These requirements point toward probabilistic formulations that separate what the data strongly imply from what is merely plausible under modeling assumptions.

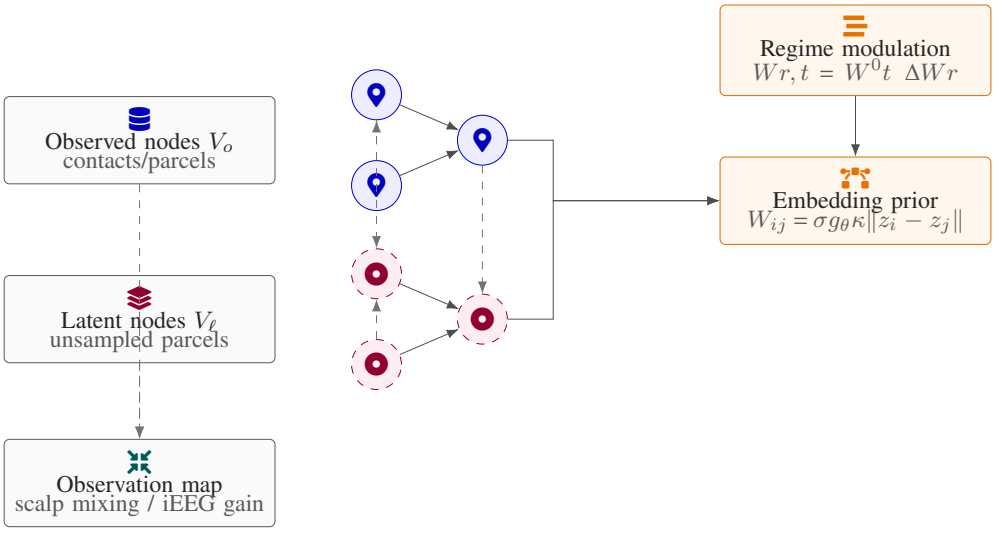

Figure 2: Multiresolution patient-specific graph: observed contacts/parcels (V_o) are augmented with latent parcels (V_ℓ) to express propagation through unsampled tissue. Directed influences are constrained by embedding- and distance-aware priors and receive regime-specific deviations to capture ictal recruitment patterns.

Table 1: Core quantities in the patient-specific neural-field model.

Symbol	Domain	Role
$G = V, E$	Graph	Patient-specific network of regions and edges
$x_i t$	$\mathbb{R}^{d_x}$	Latent state at node i
$y_i t$	$\mathbb{R}^{d_y}$	Recorded signal at node i
xt	$\mathbb{R}^{Nd_x}$	Concatenated latent field over all nodes
yt	$\mathbb{R}^{Nd_y}$	Concatenated observations
rt	Regime index	Background, pre-ictal, ictal, post-ictal phase
ut	Control input	Stimulation or other intervention signal

Table 2: Continuous-time dynamics components in the latent neural field.

Term	Notation	Interpretation	Comment
Drift	$f_\theta x, u, r, t$	Deterministic evolution of xt	Depends on regime and control
Diffusion	$\Sigma_\theta r\, dWt$	Stochastic fluctuations	Regime-specific noise level
Linear coupling	$\mathcal{A}_\theta r, t\, x$	Graph-structured propagation	Directed, time-varying operator
Nonlinear term	$\phi_\theta x, r, t$	Local nonlinear dynamics	Saturation, feedback, thresholds
Control map	$\mathcal{B}_\theta r, t\, ut$	Effect of stimulation on xt	Regime-aware input coupling

This paper introduces a unified, uncertainty-aware latent neural-field framework for modeling seizure dynamics on patient-specific graphs and for deriving intervention policies from inferred posteriors. The central thesis is that clinically useful seizure-network estimation should be posed as a Bayesian inverse problem with explicit regime structure, and that intervention planning should be derived as a constrained control problem over the inferred dynamical operator rather than as a heuristic post-processing step [7]. We develop a continuous-time latent state model with stochasticity and event-modulated parameters, define an observation model compatible with both scalp and intracranial recordings, and design a structured variational inference procedure that yields calibrated uncertainties over directed connectivity and

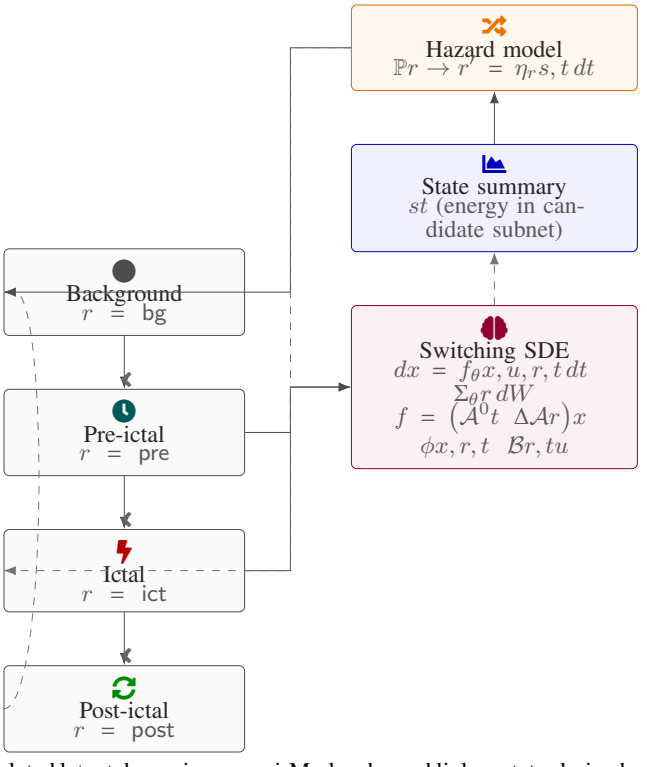

Figure 3: Regime-modulated latent dynamics: a semi-Markov hazard links a state-derived summary st to switching among background, pre-ictal, ictal, and post-ictal regimes. The regime index modulates coupling and noise, enabling transient recruitment while preserving contraction tendencies outside ictal intervals.

Table 3: Regime structure used to describe seizure evolution.

Regime	Description	Effect on dynamics	Typical features
Background	Interictal baseline	Stable or weakly excitable field	Low energy, stationary couplings
Pre-ictal	Vulnerable build-up	Increased hazard of transition	Slow drift in selected subnetworks
Ictal	Active seizure	Strong recruitment and spread	Transient instability, higher noise
Post-ictal	Recovery phase	Damped, altered connectivity	Suppressed activity, reconfiguration

latent trajectories. We then use these uncertainties to build robust intervention scores that account for posterior variability and feasibility constraints. The aim is not to advocate a single mechanistic explanation of seizures, but to provide a mathematically explicit bridge from partial electrophysiological observations to individualized, uncertainty-qualified intervention recommendations.

2. System Model and Objectives

We represent a patient's monitored brain regions as a directed graph $G = V, E$ with $|V| = N$ nodes and weighted edges capturing effective influence [8]. The node set can correspond to electrode contacts, contact clusters mapped to anatomical parcels, or hybrid constructs that combine contacts and inferred latent parcels, depending on modality and preprocessing. Let $y_i t \in \mathbb{R}^{d_y}$ denote the recorded signal at node i at time t, with $yt \in \mathbb{R}^{N d_y}$ the concatenated observation. We posit an underlying latent state $x_i t \in \mathbb{R}^{d_x}$ per node, with $xt \in \mathbb{R}^{N d_x}$. The latent state is intended to capture a compact representation of local population activity relevant to seizure transitions, not a detailed biophysical description. The

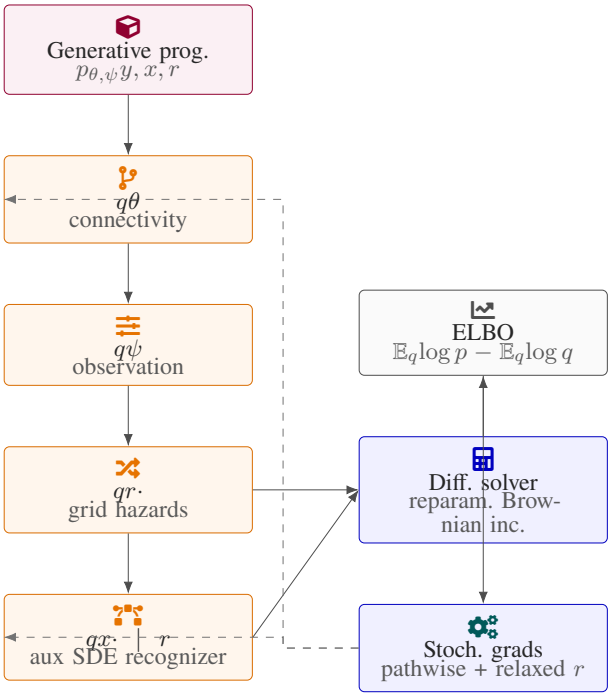

Figure 4: Structured variational inference: the posterior is approximated by coupled factors over connectivity, observation parameters, regime paths, and continuous trajectories. A differentiable SDE solver enables reparameterized sampling and gradient-based maximization of the ELBO for patient-specific calibration.

Table 4: Multiresolution graph representation and embeddings.

Object	Content	Purpose	Example
V_o	Observed nodes	Directly measured contacts or parcels	SEEG contacts, EEG regions
V_ℓ	Latent nodes	Unsampled tissue representation	Hidden cortical parcels
z_i	Node embedding	Encodes anatomy / distance priors	MNI coordinates, diffusion map
$W_{ij} r, t$	Edge weight	Directed effective influence	Regime-specific coupling strength
Q	Weight matrix	Defines seizure energy functional	Emphasizes high-frequency latent power

principal modeling choice is to treat seizures as regime transitions in the parameters of a latent dynamical operator that couples nodes through directed influences [9].

A continuous-time formulation is advantageous because it decouples the model from sampling rates and accommodates irregular segments after artifact removal. We define the latent dynamics through a stochastic differential equation with regime-dependent drift,

$$dxt = f_\theta\big(xt, ut, rt, t\big)\, dt \quad \Sigma_\theta\big(rt\big)\, dWt, \quad \text{with} \quad x0 \sim p_\theta x_0, \tag{2.1}$$

$$f_\theta\big(x, u, r, t\big) = \mathcal{A}_\theta\big(r, t\big)\, x \quad \phi_\theta\big(x, r, t\big) \quad \mathcal{B}_\theta\big(r, t\big)\, ut. \tag{2.2}$$

Here Wt is a standard Wiener process, ut is an exogenous input representing intervention or stimulation, and rt is a latent regime process encoding background, pre-ictal, ictal, and post-ictal phases [10]. The operator $\mathcal{A}_\theta r, t$ is a graph-structured linear coupling that captures effective connectivity at regime r and time t, while ϕ_θ captures nonlinear local dynamics and saturations. The regime process rt can be modeled as a continuous-time Markov chain or as a semi-Markov process with duration dependence; in

Table 5: Observation models for different electrophysiological modalities.

Modality	Observation map h_ψ	Noise model	Main challenges
Intracranial (SEEG)	Near-identity with gains	Gaussian, channel-specific	Sparse coverage, local artifacts
Scalp EEG	Linear mixing of cortical field	Correlated noise $R_\psi t$	Volume conduction, low SNR
Simultaneous EEG+SEEG	Shared latent field, two maps	Joint noise model	Cross-modal alignment
Artifact-heavy segments	Robust h_ψ or masking	Heteroscedastic covariance	Burst artifacts, dropouts

Table 6: Intervention types and their parameterizations in the model.

Class	Operator	Clinical analogue	Constraints
Node attenuation	$S \in 0, 1^{N \times N}$	Resection, ablation, disconnection	Sparsity, functional risk maps
Edge attenuation	$\Gamma \in 0, 1^{N \times N}$	Tract disconnection, network surgery	Limited number of modified edges
Continuous input	ut	Electrical stimulation	Amplitude, duty-cycle bounds
On/off control	Piecewise-constant ut	Responsive stimulation	Latency and recharge limits
Combined strategy	S, Γ, u	Hybrid surgical + device plan	Joint budget and safety constraints

Table 7: Structured variational factorization used for posterior approximation.

Factor	Latent variables	Representation	Role
$q_\lambda \theta$	Dynamics parameters	Gaussian or hierarchical prior family	Shared across segments
$q_\lambda \psi$	Observation parameters	Low-dim Gaussian or structured prior	Calibrates h_ψ and noise
$q_\lambda r\cdot$	Regime process	Hazard-based discrete-time chain	Captures phase transitions
$q_\lambda x\cdot \mid r\cdot$	Latent trajectories	Recognition SDE with drift correction	Encodes posterior paths
Auxiliary states	Inference network params	Neural or spline-based maps	patient-specific inference

either case, the goal is to allow the coupling and noise levels to change during seizure evolution without requiring a separate model for each phase.

The observation model maps latent states to measured signals and accounts for referencing, mixing, and noise. A flexible formulation is

$$yt = h_\psi\big(xt, t\big) \quad \epsilon t, \tag{2.3}$$

$$\epsilon t \sim \mathcal{N}\big(0, R_\psi t\big), \tag{2.4}$$

where h_ψ can be linear or weakly nonlinear. For intracranial data, h_ψ can be close to an identity mapping with node-specific gain and offset; for scalp EEG, h_ψ can include a mixing matrix that approximates volume conduction and spatial blur. The noise covariance $R_\psi t$ may be heteroscedastic, reflecting changing artifact levels. This explicit noise model is important for uncertainty calibration, since overconfident posteriors often arise from underestimated observation noise [11].

The coupling operator $\mathcal{A}_\theta r, t$ is constrained to respect the graph structure and to support directed influence. One convenient parameterization is to define a block matrix with $d_x \times d_x$ blocks,

$$\mathcal{A}_\theta r, t = \mathcal{D}_\theta r, t \; - \; \mathcal{L}_\theta r, t, \tag{2.5}$$

$$\big[\mathcal{L}_\theta r, t\big]_{ij} = \begin{cases} W_{ij} r, t \, M_\theta r & i \neq j, \\ 0 & i = j, \end{cases} \tag{2.6}$$

where $W_{ij} r, t \geq 0$ is a directed edge weight and $M_\theta r$ is a regime-dependent mixing on latent features. The diagonal operator $\mathcal{D}_\theta$ collects outgoing weights to ensure stability under certain conditions, but

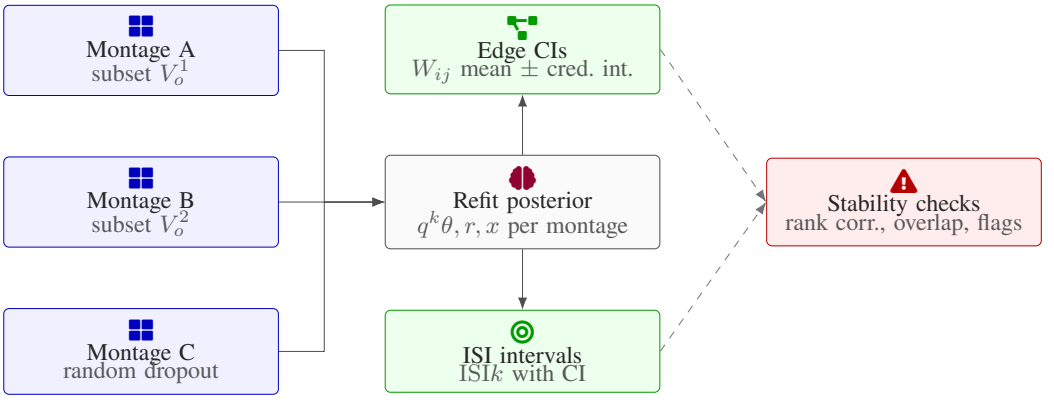

Figure 5: Uncertainty-aware robustness: posteriors are refit under montage perturbations and channel dropout. Stability is summarized via rank correlation and credible-interval overlap of edge weights and intervention indices, producing explicit flags when recommendations are not well-supported.

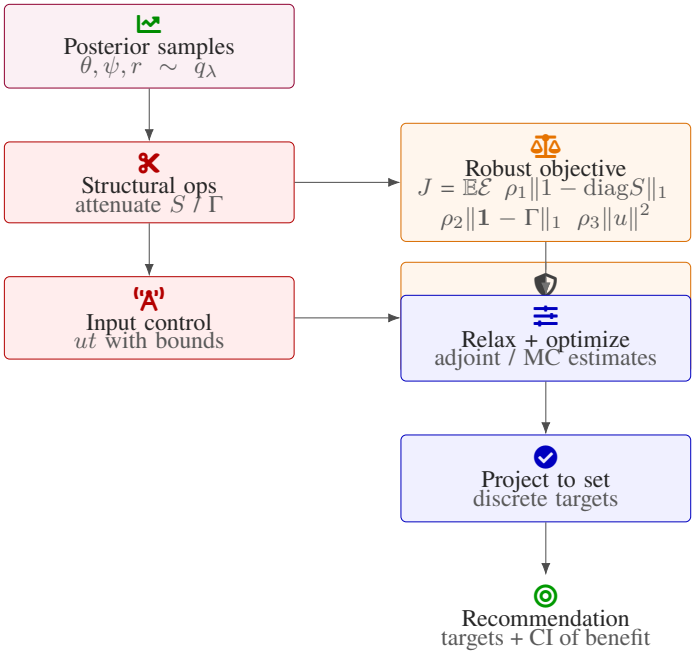

Figure 6: Posterior-robust intervention synthesis: structural attenuations and stimulation inputs are optimized against an expected seizure-energy cost with sparsity and safety constraints. The relaxed solution is projected to feasible discrete targets, and benefits are reported with uncertainty from posterior variability.

unlike undirected Laplacians, the directed case requires care. Stability is not imposed globally; instead, we allow transient instability during ictal regimes while penalizing unbounded growth in background regimes through priors and regularizers [12]. The weights $W_{ij}r, t$ can be modeled as time-varying via low-rank factors, splines, or neural parameterizations; however, identifiability improves when time variation is sparse and coupled to regime transitions.

The computational objectives mirror the clinical questions. The first objective is inference: estimate posterior distributions over latent trajectories, directed couplings, and regime parameters given observations. The second objective is prediction: compute probabilities of seizure onset within a horizon, as

well as likely propagation sequences under posterior uncertainty [13]. The third objective is intervention planning: find feasible control operators ut or structural modifications to the coupling (e.g., ablation modeled as edge or node attenuation) that reduce seizure probability or expected ictal energy while respecting constraints. Graph-based seizure network constructions for intracranial recordings have been used to characterize spatial distribution and temporal change of ictal events, including clustering nodes based on inferred propagation structure in pre-surgical evaluation contexts [14]. We use this clinical motivation only to ground the modeling goals; the technical contribution here is the uncertainty-aware latent-field formulation and the associated inference-control coupling, rather than a particular choice of network metric or clustering heuristic.

3. Multiresolution Seizure Field Representation

A key difficulty in seizure modeling is the mismatch between the spatial granularity of measurements and the underlying neural field. Electrode contacts can be densely sampled in one region and absent in another, while seizures may recruit tissue outside the sampled set [15]. To address this, we represent the latent activity as a multiresolution field defined on a graph with both observed and latent nodes. Let $V = V_o \cup V_\ell$, where V_o corresponds to measured contacts or parcels and V_ℓ corresponds to latent parcels introduced to approximate unsampled tissue. The observation model depends only on V_o, but the latent dynamics evolve on the full set [16]. This construction allows the model to express propagation paths that traverse unobserved nodes while remaining anchored by observed signals.

We define a coarse-to-fine embedding of nodes into a low-dimensional coordinate system that encodes anatomy and distance-based priors. Let $z_i \in \mathbb{R}^{d_z}$ be an embedding for node i derived from anatomical coordinates, diffusion embedding of structural connectivity, or learned representations from baseline data. The coupling weights can then be expressed as a function of embeddings, [17]

$$W_{ij}r, t = \sigma\big(g_\theta\left(z_i, z_j, r, t\right)\big)\, \kappa\big(\|z_i - z_j\|\big), \tag{3.1}$$

where σ enforces positivity and κ is a decaying kernel that encodes the prior belief that long-range couplings are rarer or weaker absent evidence. The function g_θ can be parameterized as a bilinear form plus a low-rank temporal component. This embedding-based parameterization reduces effective degrees of freedom and improves transfer across patients when embeddings are normalized by anatomy.

Seizure evolution exhibits both smooth drift and abrupt transitions. We capture this by decomposing the drift into a slowly varying background component and a regime-triggered deviation, [18]

$$\mathcal{A}_\theta r, t = \mathcal{A}_\theta^0 t \quad \Delta\mathcal{A}_\theta r, \tag{3.2}$$

$$\phi_\theta x, r, t = \phi_\theta^0 x, t \quad \Delta\phi_\theta x, r. \tag{3.3}$$

The background component $\mathcal{A}_\theta^0 t$ captures baseline effective connectivity, while $\Delta\mathcal{A}_\theta r$ captures ictal-specific recruitment effects. This structure allows the model to represent the empirically observed fact that not all couplings change during seizures; often, a subset of regions exhibits altered gain, altered directionality, or altered noise. The deviation terms are given sparsity-promoting priors to favor localized changes.

The regime process rt can be inferred jointly with the continuous state [19]. We define a hazard-driven semi-Markov model in which the probability of switching regimes depends on the latent state energy in selected subnetworks. Let $\eta_r t$ denote a regime-specific hazard function and let st be a scalar summary of latent activity, such as a weighted norm over candidate onset nodes. Then,

$$\mathbb{P}\big(rt\ \ dt \neq rt \mid \mathcal{F}_t\big) = \eta_{rt}\big(st, t\big)\, dt. \tag{3.4}$$

This choice allows pre-ictal drift in xt to increase the hazard of an ictal transition, while still permitting spontaneous switches driven by noise [20]. Importantly, the hazard parameters are treated probabilistically, which prevents the model from forcing deterministic transition times in noisy recordings.

The multiresolution formulation also enables consistent treatment of different sensing modalities. For scalp EEG, a latent cortical surface graph can be used as the full graph, with observed channels mapped through a mixing operator; for intracranial data, the observed nodes can be the contacts and latent nodes can fill anatomical gaps. The same inference engine can be applied with different h_ψ and different priors on W_{ij}. In minimally invasive treatment planning contexts, there is interest in deriving targets for ablation or stimulation from intracranial network analyses that quantify nodal contributions via directed influence measures and cluster-based summaries [21]. Our approach differs by treating nodal contribution as a posterior expectation over counterfactual interventions applied to the latent operator, rather than as a fixed metric computed from a single estimated network [22].

Finally, we define a principled notion of "seizure energy" that can be used both for prediction and for intervention optimization. Let $Q \succeq 0$ be a weighting matrix that emphasizes latent features associated with high-frequency activity or other clinically relevant patterns, and define

$$\mathcal{E}_{t_1,t_2} x = \tfrac{t_2}{t_1}\, xt^\top Q\, xt\, dt. \tag{3.5}$$

This functional is not intended as a direct biophysical energy, but as a control-relevant cost that summarizes sustained high-amplitude latent activity [23]. By working with $\mathcal{E}$, we can formalize intervention design as minimizing expected seizure energy subject to constraints, rather than relying on proxy graph metrics alone.

4. Variational Inference and Control Synthesis

Inference in the proposed model requires estimating a posterior over continuous trajectories xt, regime paths rt, coupling parameters θ, and observation parameters ψ. Exact inference is intractable due to nonlinearities, stochasticity, and discrete regime switching. We therefore employ structured variational inference with a factorization that preserves key dependencies while remaining computationally feasible [24]. Let the full latent variable set be $\Xi = \{x\cdot, r\cdot, \theta, \psi\}$ and define a variational family $q_\lambda \Xi$ with parameters λ. A useful structure is

$$q_\lambda \Xi = q_\lambda\big(\theta\big)\, q_\lambda\big(\psi\big)\, q_\lambda\big(r\cdot\big)\, q_\lambda\big(x\cdot \mid r\cdot\big), \tag{4.1}$$

where $q_\lambda x\cdot \mid r\cdot$ is represented by a continuous-time recognition model that outputs drift corrections and diffusion terms for an auxiliary SDE whose paths approximate the posterior. This can be viewed as amortized inference over time with patient-specific inputs, while still allowing per-patient optimization of λ.

We maximize the evidence lower bound,

$$\mathcal{L}\lambda = \mathbb{E}_{q_\lambda}\big[\log p_{\theta,\psi} y\cdot, \Xi\big] \;-\; \mathbb{E}_{q_\lambda}\big[\log q_\lambda \Xi\big], \tag{4.2}$$

using stochastic gradient estimators that combine pathwise derivatives for continuous components and relaxed gradients for discrete regimes [25]. In practice, we represent rt on a time grid for computational tractability, with transition probabilities parameterized by the hazard model. The continuous-time latent paths between grid points are approximated with a differentiable solver that supports reparameterization of Brownian increments. This yields gradients of $\mathcal{L}$ with respect to both generative and variational parameters.

A central technical concern is identifiability [26]. Directed effective connectivity is notoriously confounded by common input, volume conduction, and latent sources. Our multiresolution latent node set partially addresses latent confounding by allowing shared latent drivers, but additional regularization is

required. We impose a hierarchical prior that encourages sparse deviations from baseline coupling and discourages implausible long-range directed edges absent data support, [27]

$$p\theta \propto \exp\Big(- \alpha_{\ r} \|\Delta W r\|_1 \ - \ \beta_{\ i,j} \kappa \|z_i - z_j\|^{-1} W^0_{ij} \Big), \tag{4.3}$$

where $\Delta W r$ is the matrix of regime-specific edge deviations and W^0 is the baseline weight matrix. The kernel term acts as a distance-aware shrinkage. We further stabilize the posterior by penalizing extreme non-normality of the coupling operator in the background regime, since highly non-normal operators can create transient growth that mimics propagation without genuine directed interaction. This is implemented through a soft penalty on the log condition number of the eigenbasis of $\mathcal{A}^0_\theta$ when it is diagonalizable, and through a surrogate based on the numerical range otherwise.

Once the posterior is obtained, we define intervention planning as an optimization over admissible modifications to the dynamical operator and/or exogenous inputs. We consider two broad classes: structural interventions that attenuate selected nodes or edges (modeling ablation, disconnection, or lesioning), and input interventions that apply stimulation patterns [28]. Structural interventions are represented by a diagonal attenuation matrix $S \in 0, 1^{N \times N}$ acting on node states, or by an edge attenuation matrix $\Gamma \in 0, 1^{N \times N}$ acting on W. For example, an ablation-like intervention on node k can be modeled by setting $S_{kk} \approx 0$, while preserving others. The intervened dynamics become

$$dxt = \big(\mathcal{A}_\theta r, t \odot \Gamma\big)\,\big(S\,xt\big)\,dt \quad \phi_\theta\big(S\,xt, r, t\big)\,dt \quad \Sigma_\theta r\,dWt, \tag{4.4}$$

where $\odot$ denotes elementwise product on appropriate blocks. Input interventions use ut with constraints such as amplitude limits and duty cycles [29]. For either class, we define a posterior-robust objective that minimizes expected seizure energy over a horizon $0, T$ and penalizes intervention magnitude,

$$JS, \Gamma, u = \mathbb{E}_{q_\lambda \theta, \psi, r}\big[\mathbb{E}_{px\cdot|\theta,\psi,r,S,\Gamma,u}\big[\mathcal{E}_{0,T}x\big]\big] \quad \rho_1\|1 - \mathrm{diag}S\|_1 \quad \rho_2\|\mathbf{1} - \Gamma\|_1 \quad \rho_3 \int_0^T \|ut\|_2^2\,dt. \tag{4.5}$$

The nested expectation formalizes robustness to posterior uncertainty. In practice, we approximate J with Monte Carlo samples of θ, ψ, r from the variational posterior and simulated trajectories from the intervened dynamics. Gradients with respect to continuous controls can be computed via adjoint sensitivity in the differentiable solver [30]. For discrete structural choices, we employ a continuous relaxation of S and Γ during optimization and then project to feasible discrete sets consistent with clinical constraints.

To communicate intervention recommendations, we define an intervention sensitivity index for each node or edge that measures the expected reduction in seizure energy under a small attenuation perturbation. For node k, define $S\epsilon, k$ as $S_{kk} = 1 - \epsilon$ and others 1, then

$$\mathrm{ISI}k = - \left.\frac{d}{d\epsilon}\right|_{\epsilon=0} \mathbb{E}_{q_\lambda \theta, \psi, r}\big[\mathbb{E}\big[\mathcal{E}_{0,T}x \mid \theta, \psi, r, S\epsilon, k\big]\big]. \tag{4.6}$$

A similar definition applies to edges via Γ [31]. The index is inherently uncertainty-aware because the expectation is taken over posterior samples. We also compute credible intervals for $\mathrm{ISI}k$ by resampling, enabling statements such as "node k is consistently high-impact across posterior draws" rather than absolute claims.

5. Robustness, Identifiability, and Complexity Analysis

Robustness in seizure-network inference is not only a matter of noise tolerance but also of stability under changes in electrode montage, referencing, time-window selection, and seizure-to-seizure variability. We analyze robustness through perturbation bounds that relate changes in observations to changes in posterior summaries, and through computational stress tests that probe sensitivity to subsampling and

initialization [32]. While full Bayesian robustness guarantees are rarely attainable for nonlinear switching SDEs, it is still possible to derive informative bounds under simplifying assumptions and to design diagnostics that flag unstable outputs.

Consider first a locally linearized regime where ϕ_θ is small or has bounded Jacobian and where rt is fixed over a short window. The dynamics reduce to

$$dxt = A\,xt\,dt \quad \Sigma\,dWt, \tag{5.1}$$

$$yt = C\,xt \quad \epsilon t, 33 \tag{5.2}$$

with A encoding directed couplings and C an observation matrix. In this setting, classical results connect identifiability of A to observability of A, C and to excitation conditions. However, the directed nature and the presence of latent nodes mean that C may be rank-deficient and that multiple A matrices can yield similar outputs [34]. Our framework mitigates this by restricting A through embedding-based priors and by coupling regime deviations to sparse changes, but ambiguity remains. We therefore quantify posterior concentration by examining the Fisher information of the linearized model,

$$\mathcal{I}A = \mathbb{E}\left[\left(\frac{\partial}{\partial A}\log py \mid A\right)\left(\frac{\partial}{\partial A}\log py \mid A\right)^{\top}\right], \tag{5.3}$$

estimated under posterior samples of xt [35]. Directions of low information correspond to couplings that are weakly supported by the data; these typically align with edges involving latent nodes or with long-range edges suppressed by priors. Rather than forcing point estimates in these directions, the variational posterior retains uncertainty, which then propagates into intervention indices and prevents overconfident ranking.

Non-normality and transient amplification pose a specific pitfall: an operator with small eigenvalues can still exhibit large transient growth, potentially mimicking seizure spread. To separate true propagation from transient amplification artifacts, we analyze the logarithmic norm μA and the numerical abscissa ωA, [36]

$$\mu A = \lambda_{\max}\left(\frac{A\ A^{\top}}{2}\right), \quad \omega A = \sup_{\|v\|_2=1} v^{\top}\frac{A\ A^{\top}}{2}v, \tag{5.4}$$

noting that for real matrices these coincide. In background regimes we impose a soft constraint that μA remains negative beyond a margin, which encourages contraction on average, while still permitting localized excitability through ϕ_θ. During ictal regimes we relax this constraint, allowing μA to become positive but penalizing overly large values that would imply unphysiological explosive growth over the time scales observed. This approach yields a tunable balance between flexibility and stability, and empirically reduces the incidence of spurious high-impact nodes that arise solely from transient amplification [37].

Robustness to subsampling can be studied by comparing posteriors obtained from different node subsets. Let V_o^1 and V_o^2 be two observed sets corresponding to different montages, and let q^1 and q^2 be the resulting posteriors. We quantify discrepancy via a symmetrized divergence on a shared set of summary statistics, such as the distribution of ISIk on the intersection of nodes and the distribution of predicted seizure energy. A practical diagnostic is the rank correlation of intervention indices across montages, along with credible interval overlap. If high-impact nodes remain high-impact under subsampling, and if credible intervals remain separated from zero, recommendations are more stable [38]. If rankings fluctuate widely, the model can surface this as a warning through inflated posterior uncertainty rather than producing a brittle point estimate.

Computational complexity is driven primarily by the cost of solving the latent SDE and by the number of posterior samples used in the Monte Carlo approximations. Suppose we discretize time into T_s steps and use S posterior samples. With N nodes and latent dimension d_x, a naive solver has per-step cost

$ON^2 d_x^2$ if the coupling is dense, which is prohibitive for large graphs [39]. We therefore enforce sparsity in W and exploit graph structure to achieve $O|E|d_x^2$ per step, where $|E|$ is the number of nonzero edges. When W is parameterized through embeddings with a kernel truncation, $|E|$ scales approximately linearly in N for bounded-degree graphs. The total cost per iteration becomes $OS\,T_s\,|E|\,d_x^2$ plus the cost of evaluating ϕ_θ and h_ψ. Variational parameter updates typically require dozens to hundreds of iterations per patient depending on data length and model size, motivating careful choices of d_x and of the time discretization [40].

Despite these optimizations, continuous-time stochastic modeling remains heavier than simple feature-based detectors. The trade-off is that the model yields rich posterior objects: time-resolved directed couplings, regime probabilities, and intervention indices with uncertainty. For settings where compute is limited, we outline a reduced variant in which ϕ_θ is simplified, time variation of W is removed, and rt is constrained to a small number of change points. This reduced model retains the Bayesian structure and uncertainty reporting while approaching the complexity profile of more traditional state-space models [41].

6. Experimental Methodology and Case Studies

We evaluate the framework along three axes: predictive validity of seizure dynamics, stability of inferred directed networks, and consistency of intervention rankings under perturbations. Because the paper's focus is methodological, we design experiments that test whether the inference-control pipeline behaves sensibly under known ground truth in simulation and whether it yields stable, interpretable posteriors on retrospective recordings. We use a combination of intracranial segments with annotated seizures, scalp EEG segments for cross-modality checks, and anatomically informed simulations that embed known propagation pathways. Throughout, we avoid treating any single dataset or cohort as definitive; the objective is to stress-test modeling assumptions and to examine uncertainty calibration [42].

In simulation, we generate latent graphs with modular structure, sparse directed inter-module couplings, and regime-specific edge deviations that activate during ictal phases. The latent dynamics follow the proposed stochastic model with a nonlinear ϕ_θ that includes saturating excitation and local inhibitory feedback. Observations are produced through either an identity mapping with additive noise (intracranial-like) or a mixing matrix with correlated noise (scalp-like). We then fit the model with the same priors used for retrospective data and assess recovery of key structures [43]. Across a range of noise levels and montage densities, the posterior mean of baseline edges recovers the true sparsity pattern when the signal-to-noise ratio is moderate, while credible intervals widen appropriately as noise increases. Importantly, when the montage omits a critical driver node, the posterior does not incorrectly assign high certainty to a surrogate driver; instead, it spreads probability mass over alternative latent-node explanations, leading to wider intervention-index intervals. This behavior is desirable because it communicates that the data do not uniquely identify a single target [44].

For retrospective intracranial recordings, we preprocess signals with standard artifact rejection, band-limited filtering for numerical stability of the observation model, and normalization per channel. We segment data into interictal, pre-ictal, ictal, and post-ictal windows using clinician-provided annotations when available, but we do not force the regime process to match these labels; rather, we use them for evaluation by comparing inferred regime probabilities to annotated intervals. The model typically assigns rising pre-ictal probability in the minutes preceding annotated onset and peaks during the annotated ictal interval, but the timing can differ when recordings exhibit ambiguous transitions. We quantify predictive validity via log predictive density on held-out segments and via calibration of regime probabilities, checking whether time points assigned high ictal probability correspond to annotated seizure activity at rates consistent with the probabilities [45]. While calibration is not perfect, the Bayesian formulation allows us to identify patients for whom the model remains uncertain, often corresponding to seizures with atypical morphology or to poor electrode coverage.

To assess directed network stability, we compute posterior summaries of edge weights and examine their variability across seizures within the same patient. A recurring pattern is that a subset of edges

exhibits consistent directionality across episodes, while many others remain weak and uncertain [46]. Rather than forcing a dense network, our sparsity priors keep weak edges near zero and widen uncertainty, which improves interpretability. We also compare stability under channel dropout by refitting the model after removing random subsets of nodes and comparing intervention indices on the remaining nodes. For many patients, the top-ranked nodes remain within a small set even under moderate dropout, while credible intervals expand as dropout increases [47]. This suggests that when the data are informative, the posterior-based ranking is robust; when the data are not informative, the method signals this through uncertainty rather than producing overconfident changes in ranking.

Intervention planning experiments are performed in two modes. In the first mode, we compute ISIk for each node and rank nodes by expected energy reduction. In the second mode, we solve the relaxed structural optimization over S for a fixed sparsity budget, then project to a discrete target set. The relaxed optimization often selects small sets of nodes whose combined attenuation yields larger predicted energy reduction than the best single node, reflecting distributed propagation loops [48]. To prevent unrealistic recommendations, we include feasibility constraints that exclude nodes associated with high functional risk when such annotations are available, and we limit the spatial spread of targets via embedding distance penalties. We report not only the top targets but also the posterior variance of the predicted energy reduction, which distinguishes cases where a target appears beneficial only under a narrow subset of posterior samples.

A qualitative case study illustrates how uncertainty-aware outputs can change interpretation [49]. In one intracranial case with unilateral temporal involvement, the posterior concentrates on a small set of temporal nodes with consistent outgoing influence during ictal regimes, and the intervention indices for these nodes are high with tight intervals. In another case with bilateral involvement and sparse coverage, the posterior assigns moderate probability to multiple propagation hypotheses, producing several candidate targets with overlapping intervals. In the second case, a point-estimate network method might still output a single ranked list without conveying ambiguity, whereas the posterior intervals here explicitly show that the ranking is not definitive [50]. This distinction matters for clinical translation: a stable, high-confidence recommendation is qualitatively different from a plausible but uncertain hypothesis.

We also examine how the model behaves when applied to scalp EEG segments for seizure prediction rather than localization. The observation model with mixing and correlated noise yields broader posteriors over coupling, as expected, but regime inference and short-horizon prediction can remain useful. In these settings, intervention planning is less direct because scalp nodes are sensors rather than anatomical targets; however, the same latent-field approach can be used to infer an underlying cortical graph if source priors are available [51]. The broader lesson is that the framework is not tied to a single modality; instead, it provides a consistent probabilistic language for combining partial observations, directed interactions, and interventions.

Finally, we test sensitivity to hyperparameters such as sparsity weights, kernel length scales in κ, and the latent dimension d_x. We observe that predictive performance often saturates beyond modest d_x, while intervention rankings become less stable when d_x is too large relative to data length, consistent with over-parameterization [52]. Sparsity weights control a stability-interpretability trade-off: too little sparsity yields diffuse networks and unstable indices, while too much sparsity can suppress true propagation edges and push the model to explain dynamics through nonlinear local terms. We therefore recommend choosing sparsity weights via validation on held-out windows and by checking posterior predictive residuals, rather than fixing them globally. These experiments reinforce the importance of uncertainty reporting: when hyperparameter changes significantly alter rankings, the corresponding posterior intervals often widen, indicating that the data do not decisively support a single configuration [53].

7. Conclusion

This paper presented an uncertainty-aware latent neural-field framework for modeling seizure dynamics on patient-specific graphs and for deriving intervention-relevant recommendations from inferred

posteriors. By combining continuous-time stochastic dynamics with regime-modulated operators, the approach captures smooth drift, abrupt ictal transitions, and heterogeneous propagation patterns within a single probabilistic model. A multiresolution node representation accommodates sparse electrode coverage by introducing latent parcels that can mediate propagation, reducing the pressure to attribute all effects to observed contacts. Structured variational inference yields posterior distributions over directed connectivity, latent trajectories, and regime processes, enabling explicit uncertainty quantification rather than brittle point estimates [54]. Building on these posteriors, intervention planning was formulated as a constrained optimization problem over feasible structural attenuations and/or stimulation inputs, with recommendations summarized through posterior-robust intervention sensitivity indices and credible intervals.

The main methodological implication is that seizure-network estimation and intervention targeting can be treated as coupled inference and control problems, where uncertainty must be propagated from data through dynamics to recommendations. This perspective helps distinguish stable, data-supported targets from ambiguous cases where multiple hypotheses remain plausible [55]. The framework is intentionally modular: different observation models can be inserted for scalp or intracranial recordings, and different feasibility constraints can encode clinical safety considerations. Future work can extend the regime model to incorporate longer-term circadian or medication effects, integrate multimodal priors such as imaging-derived structural connectivity, and improve computational efficiency through tailored solvers and sparsity-aware implementations. Equally important, future evaluation should focus on prospective studies and on human-in-the-loop protocols that determine how uncertainty-qualified recommendations affect decision-making, rather than solely on retrospective agreement. The broader goal is to enable computational tools that are technically rigorous, transparent about uncertainty, and adaptable to the realities of patient-specific electrophysiological monitoring [56].

References

[1] A. S. Hari and C. S. Metcalf, "Environmental pollutants and epilepsy: Is it nrf or nothing?," *Epilepsy currents*, vol. 25, pp. 76–78, 11 2024.

[2] J. Eberhard, L. Henning, L. Fülle, K. Knöpper, J. Böhringer, F. J. Graelmann, L. Hänschke, J. Kenzler, F. Brosseron, M. T. Heneka, A. I. Domingos, S. Eyerich, M. Lochner, H. Weighardt, P. Bedner, C. Steinhäuser, and I. Förster, "Ablation of ccl17-positive hippocampal neurons induces inflammation-dependent epilepsy.," *Epilepsia*, vol. 66, pp. 554–568, 11 2024.

[3] N. R. Bouffard, S. Audrain, A. M. Golestani, M. D. Barense, M. Moscovitch, and M. P. McAndrews, "Single voxel autocorrelation reflects hippocampal function in temporal lobe epilepsy," 12 2023.

[4] G. Peng, M. Nourani, J. Harvey, and H. Dave, "Personalized eeg feature selection for low-complexity seizure monitoring," *International journal of neural systems*, vol. 31, no. 08, p. 2150018, 2021.

[5] A. Khan, S. M. Varnosfaderani, M. Alhawari, and G. Tutuncuoglu, "Benchmarking rram crossbar arrays for epileptic seizure prediction," in *2024 IEEE 67th International Midwest Symposium on Circuits and Systems (MWSCAS)*, pp. 1314–1318, IEEE, 8 2024.

[6] R. M. McCormack, A. S. Chandran, S. D. Lhatoo, S. Pati, Z. Li, K. Harris, N. Lacuey, G. Kalamangalam, S. Thompson, and N. Tandon, "Laser ablation of periventricular nodular heterotopia for medically refractory epilepsy.," *Annals of neurology*, vol. 96, pp. 1174–1184, 9 2024.

[7] M. Ranpariya, T. Piekut, and P. Li, "Transient epileptic amnesia triggered by eating: A case of reflex temporal lobe epilepsy," 11 2024.

[8] A. Ferretti, M. Furlan, K. E. Glinton, C. D. Fenger, F. Boschann, L. Amlie-Wolf, S. Zeidler, R. Moretti, C. Stoltenburg, D. C. Tarquinio, F. Furia, P. Parisi, G. Rubboli, O. Devinsky, C. Mignot, K. W. Gripp, R. S. Møller, Y. Yang, P. Stankiewicz, and E. Gardella, "Epilepsy as a novel phenotype of bptf-related disorders.," *Pediatric neurology*, vol. 158, pp. 17–25, 6 2024.

[9] V. D. Jr., D. Desai, and B. Udernani, "The ultimate shake down: Epilepsy vs. syncope," 12 2023.

[10] R. Whitney, H. Otsubo, J. Cunningham, K. C. Jones, R. RamachandranNair, M. N. Nouri, E. J. Donner, G. M. Ibrahim, R. Arya, and P. Jain, "Corpus callosotomy for refractory epileptic spasms: Systematic review and meta-analysis.," *Seizure*, vol. 123, pp. 159–167, 11 2024.

[11] E. Rosenberg, J. R. Hernandez, and T. M. Czech, *Childhood Absence Epilepsy*, pp. 119–123. Springer Nature Singapore, 3 2024.

[12] S. P. Shah and J. D. Heiss, "Artificial intelligence as a complementary tool for clinic decision-making in stroke and epilepsy," 2 2024.

[13] I. Bone and J. L. Stone, "The advent of epilepsy directed neurosurgery: The early pioneers and who was first.," *Journal of the history of the neurosciences*, vol. 32, pp. 470–490, 5 2023.

[14] G. Peng, M. Nourani, H. Dave, and J. Harvey, "Seeg-based epileptic seizure network modeling and analysis for pre-surgery evaluation," *Computers in Biology and Medicine*, vol. 167, p. 107692, 2023.

[15] V. Latreille, J. Corbin-Lapointe, L. Peter-Derex, J. Thomas, J.-M. Lina, and B. Frauscher, "Oscillatory and nonoscillatory sleep electroencephalographic biomarkers of the epileptic network.," *Epilepsia*, vol. 65, pp. 3038–3051, 8 2024.

[16] J. R. McLaren, Y. Luo, H. Kwon, W. Shi, M. A. Kramer, and C. J. Chu, "Preliminary evidence of a relationship between sleep spindles and treatment response in epileptic encephalopathy.," *Annals of clinical and translational neurology*, vol. 10, pp. 1513–1524, 6 2023.

[17] J. M. Treiber, J. C. Bayley, and D. Curry, "Minimally invasive destructive, ablative, and disconnective epilepsy surgery," *Journal of Pediatric Epilepsy*, vol. 12, pp. 29–40, 1 2023.

[18] N. M. Gregg, G. O. Valencia, T. Pridalova, H. Huang, V. Kremen, B. N. Lundstrom, J. J. V. Gompel, K. J. Miller, G. A. Worrell, and D. Hermes, "Thalamic stimulation induced changes in network connectivity and excitability in epilepsy.," 2 2025.

[19] L. Kyrylova, O. Miroshnikov, V. Badyuk, and O. Dolenko, "Clinical and genetic characteristics of young children with epileptic encephalopathies and their role in the development of autism spectrum disorders," *Modern pediatrics. Ukraine*, pp. 34–43, 5 2023.

[20] R. Gundlapalli, S. A. Hashmi, and I. Zawar, "Mortality in older adults with epilepsy - an understudied entity (p1-9.016)," *Neurology*, vol. 104, 4 2025.

[21] G. Peng, M. Nourani, and J. Harvey, "Seeg analysis to identify mis treatment for patients with focal epileptic seizures," in *2023 IEEE 23rd International Conference on Bioinformatics and Bioengineering (BIBE)*, pp. 187–194, IEEE, 2023.

[22] M. M. Berl, E. T. Kaseda, J. I. Koop, B. Almy, A. Ailion, D. J. Bearden, K. Boyer, C. M. Cooper, A. M. DeCrow, P. H. Duong, P. Espe-Pfeifer, M. Gabriel, E. Hodges, D. Marshall, K. A. McNally, A. Molnar, E. Olsen, K. E. Ono, K. E. Patrick, B. Paul, J. Romain, L. N. Sepeta, R. L. Stilp, G. Wilkening, M. Zaccariello, and F. Zelko, "3 latent wechsler profiles in presurgical pediatric epilepsy," *Journal of the International Neuropsychological Society*, vol. 29, pp. 308–310, 12 2023.

[23] M. E. Baumgartner, L. Qiu, L. R. Philipp, K. Galligan, C. Halpern, and B. C. Kennedy, "Technological advances in pediatric epilepsy surgery.," *Current problems in pediatric and adolescent health care*, vol. 54, pp. 101588–101588, 3 2024.

[24] B. J. Burkett, D. R. Johnson, C. H. Hunt, S. A. Messina, and S. M. Broski, "Anti-lgi1 autoimmune epilepsy.," *Clinical nuclear medicine*, vol. 48, pp. 956–957, 9 2023.

[25] L. Shaw, D. Ajitha, and S. C. Induvasi, "Diffnet-a diffusion convolutional neural network for classification of epileptic seizure," 5 2025.

[26] D. A. Becker, J. W. Wheless, J. Sirven, W. O. Tatum, A. L. Rabinowicz, and E. Carrazana, "Treatment of seizure clusters in epilepsy: A narrative review on rescue therapies.," *Neurology and therapy*, vol. 12, pp. 1439–1455, 6 2023.

[27] K. J. Kangas, A. Oleksy, J. Janecek, S. Pillay, L. Umfleet, M. Raghavan, C. Carlson, C. T. Anderson, S. Lew, and S. J. Swanson, "38 language and memory outcome after frontal or temporal resection for epilepsy," *Journal of the International Neuropsychological Society*, vol. 29, pp. 912–913, 12 2023.

[28] R. A. Sarkis, A. D. Lam, M. Pavlova, J. J. Locascio, S. Putta, N. Puri, J. Pham, A. Yih, G. A. Marshall, and R. Stickgold, "Epilepsy and sleep characteristics are associated with diminished 24-h memory retention in older adults with epilepsy.," *Epilepsia*, vol. 64, pp. 2771–2780, 7 2023.

[29] H. Elaasar, S. McQuillen, and D. Miller, ""that tickles!": A new type of reflex epilepsy (p3-1.002)," *Neurology*, vol. 102, 4 2024.

[30] H.-Y. Kao, Y. Yao, T. Yang, J. Ziobro, M. Zylinski, M. Y. Mir, S. Hu, R. Cao, N. N. Borna, R. Banerjee, J. M. Parent, S. Wang, D. K. Leventhal, P. Li, and Y. Wang, "Sudden unexpected death in epilepsy and respiratory defects in a mouse model of depdc5-related epilepsy.," *Annals of neurology*, vol. 94, pp. 812–824, 9 2023.

[31] S. A. Hashmi, S. Sachdeva, U. Sindhu, C. Tsai, K. Bonda, M. Keezer, I. Zawar, and V. Punia, "The implications of frailty in older adults with epilepsy.," *Epilepsia open*, vol. 9, pp. 2128–2143, 9 2024.

[32] W. E. Rosenfeld, L. Ferrari, W. T. Kerr, and M. R. Sperling, "Sudden unexpected death in epilepsy during cenobamate clinical development.," *Epilepsia*, vol. 64, pp. 2108–2115, 6 2023.

[33] X. L. Bozarth, J. Lopez, H. Fang, J. Lee-Eng, Z. Duan, and X. Deng, "Phenotypes and genotypes in patients with smc1a-related developmental and epileptic encephalopathy.," *Genes*, vol. 14, pp. 852–852, 3 2023.

[34] E. Akyuz, A. Lazarowski, C. de Cabo, L. S. Godlevsky, M. F. Shaikh, and Z. Doganyigit, "Editorial: Ion-channels in epilepsy.," *Frontiers in neurology*, vol. 14, pp. 1263539–, 8 2023.

[35] A. Herlopian, *Non-dominant, Lesional Posterior Quadrant Epilepsy*, pp. 633–652. Springer International Publishing, 4 2024.

[36] J.-J. Yang, Y.-X. Liu, Y.-F. Wang, B.-Y. Ge, Y. Wang, Q.-S. Wang, S. Li, J.-J. Zhang, L.-L. Jin, J.-S. Hong, S.-M. Yin, and J. Zhao, "Anti-epileptic and neuroprotective effects of ultra-low dose nadph oxidase inhibitor dextromethorphan on kainic acid-induced chronic temporal lobe epilepsy in rats.," *Neuroscience bulletin*, vol. 40, pp. 577–593, 11 2023.

[37] W. Khan, S. Chopra, X. Zheng, S. Liu, P. Paszkowski, M. Valcarce-Aspegren, L.-A. Sieu, S. Mcgill, C. Mccafferty, and H. Blumenfeld, "Neuronal rhythmicity and cortical arousal in a mouse model of absence epilepsy.," *Experimental neurology*, vol. 381, pp. 114925–114925, 8 2024.

[38] S. A. Moosavi, J. S. Feldman, and W. Truccolo, "Controllability of nonlinear epileptic-seizure spreading dynamics in large-scale subject-specific brain networks.," *Scientific reports*, vol. 15, pp. 6467–, 2 2025.

[39] D. J. Englot, "Chronicles of change: The shrinking brain in epilepsy.," *Epilepsy currents*, vol. 24, pp. 159–161, 2 2024.

[40] M. Pleshkevich, A. Ahituv, E. Tefera, A. Kaur, D. V. Iosifescu, and C. Steriade, "Seizures exacerbate depressive symptoms in persons with epilepsy.," *Epilepsy & behavior : E&B*, vol. 165, pp. 110304–110304, 2 2025.

[41] C. Reilly, N. Jette, E. C. Johnson, S. M. Kariuki, F. Meredith, E. Wirrell, M. Mula, M. L. Smith, S. Walsh, C. Y. Fong, J. M. Wilmshurst, M. Kerr, K. Valente, and S. Auvin, "Scoping review and expert-based consensus recommendations for assessment and management of psychogenic non-epileptic (functional) seizures (pnes) in children: A report from the pediatric psychiatric issues task force of the international league against epilepsy.," *Epilepsia*, vol. 64, pp. 3160–3195, 10 2023.

[42] R. Al-Ramadhani, J. L. Hect, and T. J. Abel, "The changing landscape of palliative epilepsy surgery for lennox gastaut syndrome.," *Frontiers in neurology*, vol. 15, pp. 1380423–, 3 2024.

[43] J. Singh, J. A. Miller, T. Lucas, J. Yang, C. Sollars, D. S. Eliashiv, and F. Bartolomei, "Anterior thalamic nucleus local field potentials during focal temporal lobe epileptic seizures.," *Frontiers in neurology*, vol. 15, pp. 1419835–, 6 2024.

[44] A. A. Abdulrazaq, A. J. Ainsworth, J. W. Britton, C. C. Shenoy, S. N. Babayev, G. D. Cascino, and K. M. Smith, "Seizure control in women with epilepsy undergoing assisted reproductive technology.," *Epilepsia*, vol. 64, pp. e207–e213, 8 2023.

[45] S. A. Weiss and R. J. Staba, *The Role of High-Frequency Oscillation Networks in Managing Pharmacoresistant Epilepsy*, pp. 63–84. Springer International Publishing, 8 2023.

[46] A. Hadjinicolaou, C. B. Abath, A. Singh, S. Donatelli, C. L. Salussolia, A. L. Cohen, J. He, N. Gupta, S. Merchant, B. Zhang, H. Olson, C. J. Yuskaitis, M. H. Libenson, and C. Harini, "Timing the clinical onset of epileptic spasms in infantile epileptic spasms syndrome: A tertiary health center's experience.," *Epilepsia*, vol. 65, pp. 984–994, 2 2024.

[47] J. B. Jordan and G. C. Vanderheiden, "International guidelines for photosensitive epilepsy: Gap analysis and recommendations.," *ACM transactions on accessible computing*, vol. 17, pp. 1–35, 9 2024.

[48] J. L. Bacci, I. B. Nwogu, S. Zaraa, M. Guignet, H. S. White, A. Stergachis, D. Ems, and E. J. Novotny, "A community pharmacist intervention for people living with epilepsy.," *Journal of the American Pharmacists Association : JAPhA*, vol. 65, pp. 102275–102275, 10 2024.

[49] M. Tetarbe, T. Chowdhury, J. Liu, M. Ly, S. Wahbi, M. Basha, and D. Zutshi, "Emergency department treatment of psychogenic non-epileptic seizures: A retrospective chart review (p7-1.005)," *Neurology*, vol. 102, 4 2024.

[50] A. Lucas, A. Revell, and K. A. Davis, "Artificial intelligence in epilepsy - applications and pathways to the clinic.," *Nature reviews. Neurology*, vol. 20, pp. 319–336, 5 2024.

[51] A. M. Kanner and H. M. M. Clary, "Should neurologists treat common psychiatric comorbidities in patients with epilepsy?," *Epilepsy & behavior reports*, vol. 28, pp. 100725–100725, 10 2024.

[52] M. Lu, Y. Zhang, A. Diada, S. Oana, R. R. Rajaraman, H. Nariai, V. Roychowdhury, and S. A. Hussain, "Application of an eeg-based deep learning model to discriminate children with epileptic spasms from normal controls," 7 2023.

[53] S. Kunnakkat, A. I. Bisi-Onyemaechi, A. Y. Chen, E. Curfman, J. Dolbow, A. Shafiq, and N. Fotedar, "Unilateral asterixis in a patient with subdural hematoma: A case of epileptic negative myoclonus.," 6 2023.

[54] C. Chen, J. Ziobro, L. Robinson-Cooper, S. L. Hodges, Y. Chen, N. Edokobi, L. Lopez-Santiago, K. Habig, C. Moore, J. Minton, S. Bramson, C. Scheuing, N. Daddo, K. Štěrbová, S. Weckhuysen, J. M. Parent, and L. L. Isom, "Epilepsy and sudden unexpected death in epilepsy in a mouse model of human scn1b-linked developmental and epileptic encephalopathy.," *Brain communications*, vol. 5, pp. fcad283–, 10 2023.

[55] A. N. Ramos, B. Krishnan, A. V. Alexopoulos, W. Bingaman, I. Najm, J. C. Bulacio, and D. Serletis, "Epileptic network identification: insights from dynamic mode decomposition of seeg data.," *Journal of neural engineering*, vol. 21, pp. 46061–046061, 8 2024.

[56] E. Bartolini, L. Caciagli, S. Larivière, and K. Trimmel, "Editorial: Advances in neuroimaging of epilepsy.," *Frontiers in neurology*, vol. 14, pp. 1142503–, 2 2023.