# BAYESENDS: BAYESIAN ELECTROPHYSIOLOGICAL NEURAL DYNAMICAL SYSTEMS FOR ALZHEIMER'S DISEASE DIAGNOSIS

**Anonymous authors**Paper under double-blind review

#### **ABSTRACT**

Alzheimer's disease (AD) alters Electroencephalogram (EEG) through slowed oscillations and diminished neural drive, yet most AD-EEG pipelines are black-box classifiers, lacking a unifying mathematical account of how both neural activity and its interaction dynamics evolve over time. We introduce BayesENDS, a Bayesian electrophysiological neural dynamical system that explores the possibility of incorporating neuron spiking mechanisms into a Bayesian neural dynamical system. By introducing a differentiable leaky-integrate-and-fire (dLIF) prior, BayesENDS is capable of inferring population events and interaction dynamics directly from EEG—without spike or interaction annotations. The dLIF prior encodes membrane dynamics, rate/refractory constraints, and physiologically plausible frequency ranges, improving identifiability while yielding biologically plausible, subject-level biomarkers alongside AD predictions. Across synthetic event-sequence benchmarks and real AD EEG datasets, BayesENDS delivers superior performance to state-of-the-art baseline methods.

## 1 Introduction

Alzheimer's disease (AD) is a progressive neurodegenerative disorder with growing global impact. Electroencephalography (EEG) provides a non-invasive, low-cost window into brain function and consistently shows *oscillatory slowing* in AD—power increases in delta/theta and decreases in alpha/beta—together with alterations in large-scale interactions and synchrony (Jeong, 2004; Dauwels et al., 2010; Babiloni et al., 2021). While deep learning has advanced EEG-based AD assessment, most pipelines remain *black-box classifiers* (Ieracitano et al., 2020; Pineda et al., 2019; Vicchietti et al., 2023; Tawhid et al., 2025) optimized for accuracy from hand-crafted or learned features, offering limited insight into how *neural activity* and *interaction dynamics* co-evolve over time (Ehteshamzad et al., 2024; Acharya et al., 2025; Wang et al., 2024; Klepl et al., 2024).

Two technical obstacles motivate a unifying, electrophysiology-aware dynamical framework. First, scalp EEG is a noisy, frequency-dependent *linear mixture* of mesoscopic sources; recovering latent population activity is an ill-posed inverse problem sensitive to modeling choices (Michel & Brunet, 2019; Michel et al., 2004). Second, interaction metrics face *interpretational pitfalls* (volume conduction, common input, SNR differences) and can yield inconsistent estimates across analysis pipelines unless dynamics and biophysical constraints are handled explicitly (Bastos & Schoffelen, 2016; Mahjoory et al., 2017).

We address these gaps with **BayesENDS**, a *Bayesian electrophysiological neural dynamical system* that learns event-driven latent dynamics and a conditional interaction graph directly from multichannel EEG. Methodologically, BayesENDS (a) represents per-channel activity with an **Event Posterior Differential Equation (EPDE)** whose solution yields expected next-event times; (b) samples inter-event intervals via a **Mean–Evolving Lognormal Process (MELP)**, where the EPDE outputs parameterize the means of a log-normal mixture with reparameterized sampling; (c) imposes an electrophysiology-informed **differentiable leaky-integrate-and-fire (dLIF) prior** that encodes leak, refractory/rate constraints, and plausible frequency ranges; and (d) infers a directed **event-relational graph** (ERG) by mapping cross-channel event lags through a smooth nonlinearity into edge weights. The entire model is trained end-to-end with a compact variational objective;

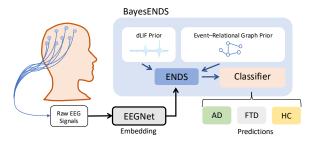


Figure 1: **Overview of the BayesENDS pipeline** illustrating the integration of raw EEG signals, EEGNet embeddings, the differentiable leaky-integrate-and-fire (dLIF) prior, and the event-relational graph (ERG) prior to generate clinically interpretable predictions for Alzheimer's Disease (AD), Frontotemporal Dementia (FTD), and Healthy Controls (HC).

our analysis provides a tractable IVP-based bound for the event-prior KL under the dLIF rate and establishes ERG stability to lag noise.

# **Major Contributions:**

- A unified Bayesian neural dynamical system that infers latent events and event–relational graph dynamics directly from EEG *without* spike or edge annotations.
- An electrophysiology-informed dLIF prior integrated into training, providing biophysical rate and refractory constraints.
- Theory establishing a computationally tractable IVP-based upper bound for the learning objective and a stability bound for BayesENDS's inferred graph dynamics.
- Empirical evidence showing: (i) accurate recovery of latent event and graph dynamics that improve understanding of AD; and (ii) superior performance over strong baselines on both synthetic benchmarks and real AD EEG datasets.

# 2 RELATED WORK

Most AD–EEG studies cast diagnosis as supervised classification over hand-crafted spectral/connectivity features or learned representations, achieving strong performance but offering limited mechanistic insight into how timing and interactions co-evolve. Recent Scientific Reports papers exemplify this trend: comprehensive pipelines comparing computational methods for AD classification, and multi-stage classification across the AD spectrum during memory-encoding versus rest (with higher accuracy during task-evoked states) (Vicchietti et al., 2023; Kim et al., 2024). Earlier machine-learning work integrates engineered EEG features (often spectral/topographic) within multimodal classifiers for dementia discrimination (Ieracitano et al., 2020). In contrast, our approach models *latent event dynamics* and a conditional interaction graph jointly, providing a generative account of how event timing and cross-channel lags give rise to predictive structure in EEG.

# 3 PROBLEM FORMULATION

We study unsupervised latent–event and relation discovery in multichannel time series for labeled sequences. Given  $\mathcal{D}=\{(X^{(n)},Y^{(n)})\}_{n=1}^N$  with  $X^{(n)}=\{x_c^{(n)}\}_{c=1}^C$  (e.g., multichannel EEG) and labels  $Y^{(n)}\in\mathcal{Y}$  (e.g., AD vs. control), and no supervision on per-channel events or inter-channel relations, the objective is to infer (i) channel-wise latent event dynamics and (ii) a (possibly time-varying) relational/graphical structure among channels. For each channel c, let  $T_c^{(n)}=\{t_{c,k}^{(n)}\}$  denote latent event times and let  $\mathbf{p}^{(n)}(t)=\{p_c^{(n)}(t)\}_{c=1}^C$  denote the corresponding posterior event-time distributions; the relational structure is represented as a graph process  $G^{(n)}(t)$  with adjacency  $A^{(n)}(t)\in\mathbb{R}^{C\times C}$ . We aim to recover  $\{\mathbf{p}^{(n)}(t)\}_{n=1}^N$  and  $P(G^{(n)}(\cdot)\mid\mathbf{p}^{(n)}(\cdot))$  from  $\{X^{(n)}\}_{n=1}^N$ ,

while using  $(\mathbf{p}^{(n)}(\cdot), G^{(n)}(\cdot))$  as inputs to a downstream predictor for  $Y^{(n)}$ ; importantly,  $Y^{(n)}$  does not supervise the latent events or relations directly. Interaction strength is modeled as a function of the temporal *co-occurrence and ordering* of inferred events—an assumption grounded in neurobiological *spike-timing-dependent plasticity* (STDP) mechanisms (near-coincident pre/post spikes modulate synaptic efficacy) (Bi & Poo, 1998; Feldman, 2012).

# 4 BAYESIAN NEURAL DYNAMICAL SYSTEM

#### 4.1 OVERVIEW

We introduce a Bayesian neural dynamical system for multichannel sequences that represents each channel with a latent *event* process and couples channels through a conditional *event-relational graph* (ERG) driven by the timing and ordering of inferred events. This design is motivated by settings where the clinical signal resides in *when* events and *how* they align across channels, for example, in EEG for Alzheimer's disease, oscillatory slowing and disrupted coordination suggest that event timing and cross-channel alignment are predictive, while neurobiological plasticity links near-coincident spikes to stronger coupling. Concretely, the dynamics of events is summarized by an **event posterior differential equation (EPDE)** whose read-out parameterizes the means of a **mean-evolving lognormal process (MELP)** for reparameterized sampling. The ERG G(t) is inferred from event co-occurrence and lags via an STDP-shaped mapping and, together with T and the observed X, is passed to a decoder  $p_{\theta}(Y \mid X, T, G)$  for downstream prediction. Training is end-to-end via variational learning that jointly optimizes EPDE and MELP; details follow in the learning subsection.

#### 4.2 Learning

We train end-to-end with a variational objective that *jointly* optimizes the Event Posterior Differential Equation (EPDE) and the Mean–Evolving Lognormal Process (MELP), while inferring an event–relational graph (ERG) without graph annotations; a weak observable-based regularizer is weighted by  $\beta \geq 0$  (details in the ERG subsection). For labeled data  $\{(X^{(n)},Y^{(n)})\}_{n=1}^N$ , with EPDE-induced  $q_{\phi}(T^{(n)}\mid X^{(n)})$ , MELP posterior  $q_{\phi}(\tau^{(n)}\mid X^{(n)})$ , and decoder  $p_{\theta}(Y^{(n)}\mid X^{(n)},T^{(n)},G^{(n)})$ , we maximize

$$\mathcal{L}_{\text{ELBO}}(\theta, \phi, \eta) = \sum_{n=1}^{N} \mathbb{E}_{q_{\phi}} \Big[ \log p_{\theta} \big( Y^{(n)} \mid X^{(n)}, T^{(n)}, G^{(n)} \big) \Big]$$
 (1)

$$-\operatorname{KL}_{T}^{(n)} - \operatorname{KL}_{\tau}^{(n)} \tag{2}$$

+ 
$$\beta \mathcal{R}_{ERG}^{(n)} + \lambda_{LIF} \mathcal{R}_{LIF}^{(n)}$$
. (3)

Here  $\mathrm{KL}_T^{(n)} := \mathrm{KL}ig(q_\phi(T^{(n)} \mid X^{(n)}) \parallel p_{\mathrm{dLIF}}(T)ig)$  compares the EPDE path law to the electrophysiology-informed event prior  $p_{\mathrm{dLIF}}(T)$ ;  $\mathrm{KL}_\tau^{(n)} := \mathrm{KL}ig(q_\phi(\tau^{(n)} \mid X^{(n)}) \parallel p_0(\tau)ig)$  measures divergence to a lognormal(-mixture) prior;  $\mathcal{R}_{\mathrm{LIF}}^{(n)}$  softly enforces leaky-integrate-and-fire consistency on learned rate proxies with weight  $\lambda_{\mathrm{LIF}} \geq 0$ .

**Challenges.** (i)  $\mathrm{KL}_T^{(n)}$  involves path measures induced by a differential equation and is intractable in closed form (the integral over an infinite-dimensional trajectory); we use a tractable integral—rate surrogate that depends on the dLIF rate r(t) (formal bound in the Theory subsection). (ii) Enforcing the LIF prior is nontrivial because the spike function in LIF is non-differentiable, which hinders direct incorporation into neural training; we therefore replace hard spikes with differentiable rate proxies and constrain them to follow dLIF laws via  $\mathcal{R}_{\mathrm{LIF}}^{(n)}$ . (iii) ERG learning lacks ground-truth edges; we include only a light, observable-based regularizer  $\mathcal{R}_{\mathrm{ERG}}^{(n)}$  to bias edges toward experimental statistics, deferring full details to the ERG subsection.

# 4.3 PRIOR: ELECTROPHYSIOLOGY-INFORMED DLIF PRIOR

We place a biophysical prior on event timing by instantiating each channel's latent events as a renewal process whose hazard is derived from a differentiable leaky-integrate-and-fire (dLIF) ab-

straction (Burkitt, 2006). Let the (rescaled) membrane potential evolve as

$$\frac{d}{dt}u_c(t) = b_c(t) - u_c(t), \qquad b_c(t) > 1,$$

where  $b_c(t)$  is an effective drive. The implied instantaneous *firing rate* is

$$r_c(t) = \left[ -\log(1 - 1/b_c(t)) \right]^{-1}, \qquad p_{\text{dLIF}}(t) = r_c(t) \exp\left( -\int_0^t r_c(s) \, ds \right),$$

and the dLIF event prior for channel c is the renewal law  $p_{\text{dLIF}}(T_c)$  with hazard  $r_c(t)$ . We parameterize  $b_c(t)$  by a bounded neural mapping from embeddings (optionally graph–aware), e.g.  $b_c(t) = 1 + \text{softplus}(g_{\xi}(z_c(t)))$ , ensuring  $b_c(t) > 1$  and thus  $r_c(t) > 0$ . Absolute/refractory effects are incorporated through a smooth gating factor  $\alpha_c(t) \in (0,1]$  (constructed from recent events), using  $r_c(t) \leftarrow \alpha_c(t) \, r_c(t)$  to suppress implausible near–back–to–back spikes.

Because the hard spike nonlinearity is non-differentiable, we regularize *rates* rather than spikes. Concretely, the learning objective includes a dLIF consistency term

$$\mathcal{R}_{\mathrm{LIF}} = \sum_{c} \int_{0}^{S} \left( \widehat{r}_{c}(t) - r_{c}(t) \right)^{2} dt,$$

where  $\hat{r}_c(t)$  is a differentiable rate proxy read from the EPDE state; this enforces LIF laws without invoking non–differentiable spike functions (Neftci et al., 2019). The variational KL between the EPDE–induced path law  $q_\phi(T_c \mid X)$  and  $p_{\text{dLIF}}(T_c)$  is intractable in general; in the Theory subsection we replace it by a tractable integral–rate bound that depends on  $r_c(t)$ , yielding a stable surrogate for training while preserving the biophysical semantics of the prior.

# 4.4 POSTERIOR: EVENT POSTERIOR DIFFERENTIAL EQUATION (EPDE)

For each channel i, let  $q_i(t \mid x_i)$  denote the density of the next event time given the observed sequence  $x_i$ . The expected next event after  $\tilde{t}_{i-1}$  is

$$\tilde{t}_i = \int_{\tilde{t}_i}^{\infty} t \, q_i(t \mid x_i) \, dt. \tag{4}$$

This induces the IVP

$$\Phi'_{i}(t) = -t q_{i}(t \mid x_{i}), \qquad \Phi_{i}(0) = \int_{0}^{\infty} t q_{i}(t \mid x_{i}) dt, \qquad \tilde{t}_{i} = \Phi_{i}(\tilde{t}_{i-1}). \tag{5}$$

Since  $\Phi_i(0)$  is intractable, we approximate the solution with a differentiable neural surrogate

$$\tilde{t}_i = f_{\theta_{\Phi}}(\tilde{t}_{i-1}, x_i), \tag{6}$$

implemented to ensure  $\tilde{t}_i > \tilde{t}_{i-1}$ . Differentiating the surrogate  $\Phi_i(t)$  gives the approximate event-time posterior used in training:

$$q_i(t \mid x_i) = -\frac{\Phi_i'(t)}{t}. \tag{7}$$

The predicted  $\tilde{t}_i$  (or its increment) will be used as the mean parameter for the Mean–Evolving Lognormal Process (MELP); details follow next.

# 4.5 SAMPLING: MEAN-EVOLVING LOGNORMAL PROCESS (MELP)

Given the previous event time  $\tilde{t}_{c,i-1}$  on channel c, the EPDE provides a K-dimensional vector of means  $\tilde{\tau}_c = (\tilde{\tau}_{c,1},\ldots,\tilde{\tau}_{c,K}) \in \mathbb{R}_+^K$  together with mixture weights  $\mathbf{w}_c = (w_{c,1},\ldots,w_{c,K}) \in \Delta^{K-1}$  and scales  $\mathbf{s}_c = (s_{c,1},\ldots,s_{c,K}) \in \mathbb{R}_+^K$ . MELP draws the inter–event interval  $\tau_{c,i}$  from a lognormal mixture whose component means are tied to the EPDE outputs:

$$p(\tau_{c,i} \mid \tilde{t}_{c,i-1}, X) = \sum_{j=1}^{K} w_{c,j} \operatorname{LogN}(\tau_{c,i}; \mu_{c,j}, s_{c,j}^{2}),$$
(8)

$$\mu_{c,j} = \log \tilde{\tau}_{c,j} - \frac{1}{2} s_{c,j}^2 \quad \Longrightarrow \quad \mathbb{E}_{\text{LogN}(\mu_{c,j}, s_{c,j}^2)}[\tau] = \tilde{\tau}_{c,j}. \tag{9}$$

Sampling is reparameterized to keep gradients pathwise:

$$k \sim \text{Cat}(\mathbf{w}_c), \qquad \varepsilon \sim \mathcal{N}(0, 1),$$
 (10)

$$\tau_{c,i} = \exp(\mu_{c,k} + s_{c,k}\varepsilon), \qquad t_{c,i} = t_{c,i-1} + \tau_{c,i}. \tag{11}$$

During training we use a differentiable variant of Eq. (10) (e.g., Gumbel-Softmax) and take hard samples at test time. MELP guarantees positivity, captures multimodal timing statistics, and provides closed-form component-wise KL terms against a lognormal(-mixture) prior used in the learning objective; the mixture expectation  $\mathbb{E}[\tau_{c,i}] = \sum_j w_{c,j} \tilde{\tau}_{c,j}$  is also available in closed form for downstream computations (e.g., ERG lags).

#### 4.6 EVENT-RELATIONAL GRAPH (ERG)

We infer a time-varying graph G(t) from temporal lags between channels, conditioned on the EPDE/MELP posterior. For channels i, j and a latent event realization  $T \sim q_{\phi}(T \mid X)$ , let  $\Delta t_{ij}(t; T)$ be the lag between their most recent events up to t. A smooth non-linear map  $\varphi_{\eta}: \mathbb{R} \to [0,1]$  converts this lag to an edge score

$$e_{ij}(t;T) = \varphi_{\eta}(\Delta \tilde{t}_{ij}(t;T)) \in [0,1]. \tag{12}$$

We approximate the conditional edge weight via Monte Carlo with S posterior samples,

$$\bar{A}_{ij}(t) = \frac{1}{S} \sum_{s=1}^{S} \varphi_{\eta} \left( \Delta \tilde{t}_{ij}^{(s)}(t) \right), \tag{13}$$

Because no edge annotations are available, we bias  $\overline{A}$  toward observable statistics computed from X (e.g., Pearson correlations of attributes of interest) via a Fisher–z alignment:

$$z_{ij}^{\text{obs}} = \operatorname{atanh}(s_{ij}), \quad \rho_{ij} = 2\bar{A}_{ij} - 1, \qquad z_{ij}^{\text{pred}} = \operatorname{atanh}(\rho_{ij}), \tag{14}$$

$$z_{ij}^{\text{obs}} = \operatorname{atanh}(s_{ij}), \quad \rho_{ij} = 2\bar{A}_{ij} - 1, \qquad z_{ij}^{\text{pred}} = \operatorname{atanh}(\rho_{ij}), \tag{14}$$

$$\mathcal{R}_{\text{ERG}} = \sum_{i < j} \left[ \frac{(z_{ij}^{\text{obs}} - z_{ij}^{\text{pred}})^2}{2\sigma_{ij}^2} + \frac{1}{2}\log\sigma_{ij}^2 \right], \qquad \sigma_{ij} = \varepsilon + \operatorname{softplus}(\operatorname{MLP}(z_i^* || z_j^*)), \tag{15}$$

where  $z_i^{\star}$  are time-pooled node features and  $\varepsilon > 0$ . This regularizer (weighted by  $\beta$  in the learning objective) encourages consistency with experimental observables while leaving edge strengths to be determined purely by data-driven event lags.

#### 4.7 Theory

**Lemma 1** (Shift–stability of IVPs (Huang et al., 2021)). Let e > 0 and  $U \subset \mathbb{R}^n$  be open. Let  $f_1, f_2: [a-2e, a) \to \mathbb{R}^n$  be continuously differentiable with  $||f_1'|| \le M$  for some M > 0. Consider

$$y_1'(t) = f_1(t), \quad y_1(a-e) = x_1, \qquad y_2'(t) = f_2(t) = f_1(t-e), \quad y_2(a-e) = x_2.$$

Then, as  $e \to 0^+$ .

$$\lim_{e \to 0^+} \left( \lim_{t \to a} \|y_1(t) - y_2(t)\| \right) \le \lim_{e \to 0^+} \|x_1 - x_2\|.$$

**Theorem 2** (IVP-based upper bound for the event-prior KL under dLIF rates). Let q(t) be a strictly positive, integrable density on [0,S]  $(0 < S \le \infty)$ . Let the electrophysiology-informed prior be

$$p_r(t) = r(t) \exp\left(-\int_0^t r(u) du\right),$$

where  $r:[0,S] \to [a,b] \subset (0,\infty)$  is measurable with  $0 < a \le b < \infty$ . Define the change of variables  $m=-e^{-t} \in [-e^{-S},-1)$  and  $M=-\log(-m)=t$ . Set

$$g(m) := -\frac{q(M)}{m M} \log \frac{q(M)}{M r(M) e^{-\int_0^M r(u) du}}, \qquad G'(m) = g(m), \quad G(-e^{-S}) = 0.$$
 (16)

Then

$$KL(q \parallel p_r) = \int_0^S q(t) \log \frac{q(t)}{r(t)e^{-\int_0^t r}} dt = \lim_{\varepsilon \downarrow 0} G(-\varepsilon), \tag{17}$$

and for any  $\varepsilon \in (0, e^{-S})$ ,

$$\mathrm{KL}(q \parallel p_r) \leq G(-\varepsilon) + \left| G(-2\varepsilon) - G(-\varepsilon) \right| =: \mathcal{U}_{\varepsilon}, \tag{18}$$

with  $\mathcal{U}_{\varepsilon} \to \mathrm{KL}(q||p_r)$  as  $\varepsilon \downarrow 0$ .

In training we evaluate the KL term using the computable bound  $U_{\varepsilon}$  from Eq. 18 (fixed small  $\varepsilon$  and an ODE solver for Eq. (16)).

We then analyze how entry-wise perturbations of lags affect the learned ERG when the edge map is exponential. For channels  $i \neq j$  and time t, let the noise-free lag be  $\Delta t_{ij}(t;T)$  and the perturbed lag be  $\widetilde{\Delta t}_{ij}(t;T) = \Delta t_{ij}(t;T) + \xi_{ij}(t;T)$ . Define the edge map  $\phi_{\alpha}(x) = \exp(-\alpha|x|) \in [0,1]$  with slope parameter  $\alpha > 0$  and the (noise-free and perturbed) edges

$$e_{ij}(t;T) = \phi_{\alpha}(\Delta t_{ij}(t;T)), \qquad \widetilde{e}_{ij}(t;T) = \phi_{\alpha}(\widetilde{\Delta t}_{ij}(t;T)).$$

The decoder uses the Monte-Carlo, time-averaged adjacencies

$$\bar{A}_{ij} = \frac{1}{MS} \sum_{m=1}^{M} \sum_{s=1}^{S} e_{ij}(t_m; T^{(s)}), \qquad \tilde{\bar{A}}_{ij} = \frac{1}{MS} \sum_{m=1}^{M} \sum_{s=1}^{S} \tilde{e}_{ij}(t_m; T^{(s)}).$$

**Theorem 3** (Entry–wise and matrix stability). *The exponential edge map is globally*  $\alpha$ –*Lipschitz: for all*  $x, y \in \mathbb{R}$ ,

$$\left|\phi_{\alpha}(x) - \phi_{\alpha}(y)\right| \le \alpha |x - y|. \tag{19}$$

Consequently, for any (i, j, t, T),

$$\left|\widetilde{e}_{ij}(t;T) - e_{ij}(t;T)\right| \le \alpha \left|\xi_{ij}(t;T)\right|. \tag{20}$$

Averaging over time and Monte-Carlo samples yields the entry-wise bound

$$\left| \widetilde{A}_{ij} - \overline{A}_{ij} \right| \leq \alpha \, \overline{|\xi_{ij}|}, \qquad \overline{|\xi_{ij}|} := \frac{1}{MS} \sum_{m=1}^{M} \sum_{s=1}^{S} |\xi_{ij}(t_m; T^{(s)})|, \tag{21}$$

and the matrix (Frobenius-norm) bound

$$\|\tilde{\bar{A}} - \bar{A}\|_{F} \leq \frac{\alpha}{MS} \sum_{m=1}^{M} \sum_{s=1}^{S} \|\Xi^{(m,s)}\|_{F}, \qquad \Xi^{(m,s)} := \left[\xi_{ij}(t_m; T^{(s)})\right]_{i \neq j}, \tag{22}$$

hence  $\|\widetilde{A} - \overline{A}\|_F \le \alpha \|\Xi\|_F$  with  $\|\Xi\|_F$  the average Frobenius norm of lag-noise matrices.

**Corollary 4** (Deterministic and probabilistic perturbation bounds). (i) (Uniformly bounded noise). If  $|\xi_{ij}(t;T)| \leq \varepsilon_{\infty}$  almost surely, then

$$|\tilde{\bar{A}}_{ij} - \bar{A}_{ij}| \le \alpha \, \varepsilon_{\infty}, \qquad |\tilde{\bar{A}} - \bar{A}||_{F} \le \alpha \, \varepsilon_{\infty} \, \sqrt{C(C-1)}.$$
 (23)

(ii) (Sub-Gaussian noise). Suppose  $\{\xi_{ij}(t_m;T^{(s)})\}_{m,s}$  are i.i.d., mean-zero, sub-Gaussian with proxy  $\sigma^2$  (i.e.,  $\mathbb{E}e^{\lambda\xi} \leq \exp(\lambda^2\sigma^2/2)$ ). Then each difference  $\Delta_{ij} := \tilde{A}_{ij} - \bar{A}_{ij}$  is sub-Gaussian with proxy  $\alpha^2\sigma^2/(MS)$  and

$$\mathbb{P}(|\Delta_{ij}| \ge \tau) \le 2\exp\left(-\frac{MS\,\tau^2}{2\alpha^2\sigma^2}\right). \tag{24}$$

*If, in addition,*  $\xi \sim \mathcal{N}(0, \sigma^2)$ *, then* 

$$\mathbb{E}[|\Delta_{ij}|] \le \alpha \,\sigma \sqrt{\frac{2}{\pi}}, \qquad \mathbb{E}[\|\widetilde{\bar{A}} - \bar{A}\|_{F}] \le \alpha \,\sigma \sqrt{\frac{2}{\pi}} \,\sqrt{C(C-1)}. \tag{25}$$

**Implication.** Small perturbations in EPDE/MELP lags translate linearly (in  $\alpha$ ) to entry—wise changes in the ERG, and averaging over samples/time further contracts the perturbation. Thus the ERG is provably stable to modest timing noise, with explicit constants controlled by the edge map slope  $\alpha$  and the noise magnitude.

## 5 EXPERIMENTS

#### 5.1 ALZHEIMER'S DISEASE EEG DATASET EXPERIMENTS

We extensively evaluated BayesENDS on two diverse EEG datasets covering Alzheimer's disease (AD), frontotemporal dementia (FTD), mild cognitive impairment (MCI), and healthy controls.

324 325

Table 1: Results on Alzheimer's EEG datasets.

332 333 334

331

334 335 336

337 338 339

345 346 347

348

350

351

352

358 359 360

361

357

366367368369

370 371 372

3	7	2
3	7	3
3	7	4
3	7	5
3	7	6

377

Dataset AD cohort A Dataset AD cohort B Model Accuracy (%) F1 (%) Accuracy (%) F1 (%) **EEGNet** 68.10 66.49 71.37 60.85 49.38 LCADNet 70.52 68.12 72.44 LSTM 70.52 68.24 77.89 61.35 **ATCNet** 64.71 60.98 71.09 50.92 **ADFormer** 69.35 65.28 82.38 63.89 69.98 80.00 62.21 LEAD 72.68 75.03 **BayesENDS** 72.69 89.82 64.87

Table 2: Ablation of spike-informed and connectome priors in **BayesENDS** on Alzheimer's EEG Datasets. Accuracy and F1 (%) reported as mean  $\pm$  s.d. across runs.

Dataset	Variant	Accuracy (%)	F1 (%)	
	No prior	$70.52 \pm 11.83$	$65.46 \pm 13.10$	
AD cohort A	dLIF prior	$73.92 \pm 9.84$	$71.41 \pm 10.72$	
AD Conort A	ERG prior	$72.75 \pm 6.63$	$70.35 \pm 7.91$	
	Dual priors	$\textbf{75.03} \pm 8.29$	$72.69 \pm 8.16$	
	No prior	$83.22 \pm 15.10$	$60.72 \pm 12.37$	
AD cohort B	dLIF prior	$87.98 \pm 8.09$	$62.95 \pm 9.24$	
AD COHOIT B	ERG prior	$86.20 \pm 9.96$	$65.63 \pm 9.17$	
	<b>Dual priors</b>	$89.82 \pm 8.39$	$64.87 \pm 10.67$	

Across both datasets, BayesENDS outperformed state-of-the-art baselines such as CNNs, RNNs, and transformers. In Dataset AD cohort A, BayesENDS showed a clear ability to distinguish AD, FTD, and healthy participants, with consistently lower performance variability across runs. This stability highlights its effectiveness in capturing subtle neural dynamics despite EEG heterogeneity. In the more unbalanced Dataset AD cohort B, where distinctions between moderate AD, MCI, and healthy controls are subtler, BayesENDS still achieved the highest diagnostic accuracy. While some baselines showed significant performance fluctuations, BayesENDS remained robust, balancing sensitivity and specificity—demonstrating resilience to noise and distribution shifts common in real-world EEG data. Overall, BayesENDS effectively extracts clinically meaningful biomarkers from EEG, confirming its practical potential for accurate, interpretable AD diagnostics.

#### 5.2 ABLATION STUDIES

We ablated BayesENDS to assess each prior's impact on accuracy and interpretability (Table 2). On AD Cohort A, the no-prior baseline is moderate; adding the dLIF (spiking) prior boosts accuracy and AUC, while the ERG (graph) prior improves accuracy, F1, and AUC via inter-channel modeling. Combining both yields the best scores. On AD Cohort B, the pattern holds with larger gains: the spiking prior gives the biggest accuracy lift under subtler classes, and the graph prior raises F1/AUC. Their combination again performs best. Overall, each prior helps, and together they provide robust, interpretable predictions on realistic EEG.

# 6 VISUALIZATION AND INTERPRETABILITY ANALYSIS

#### 6.1 BOUNDARY TIME PREDICTION VISUALIZATION

Figure 4 visually compares the predicted versus ground-truth boundary times for STRODE and BayesENDS across three distinct frequency bands ([5–10 Hz], [10–15 Hz], and [15–20 Hz]). STRODE demonstrates noticeable deviations from the ideal diagonal alignment, suggesting challenges in accurately recovering true boundary timings, particularly at higher frequency bands. In contrast, BayesENDS consistently maintains a tighter diagonal alignment across all frequency ranges, indicating superior accuracy and robustness in capturing underlying temporal structures.

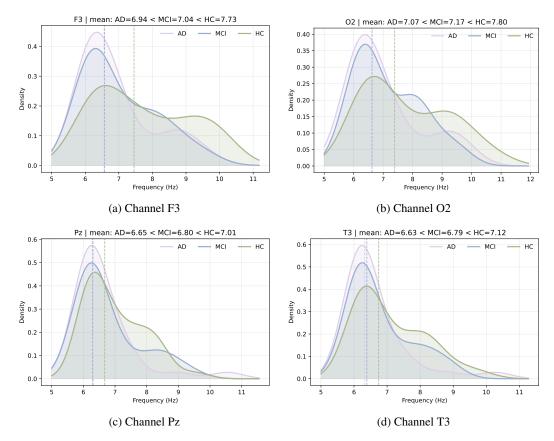


Figure 2: Kernel density estimates of the inferred dLIF frequency distributions across Alzheimer's disease (AD), mild cognitive impairment (MCI), and healthy control (HC) groups for EEG channels F3, O2, Pz, and T3. The clear trend of decreasing central frequency with increasing disease severity illustrates the physiological relevance and interpretability of the BayesENDS model's inferred latent neural dynamics.

This visualization clearly illustrates BayesENDS' effectiveness in accurately inferring latent event boundaries in the toy dataset, reinforcing its suitability for precise temporal modeling in EEG data.

#### 6.2 DLIF INFERRED FREQUENCY VISUALIZATION

To further explore and validate the interpretability of the latent neural dynamics captured by BayesENDS, we visualized the inferred frequency distributions derived from the differentiable leaky-integrate-and-fire (dLIF) prior across several key EEG channels associated with Alzheimer's disease progression, including channels F3, O2, Pz, and T3. Figure 2a exemplifies the frequency distributions at channel F3 for the Alzheimer's disease (AD), mild cognitive impairment (MCI), and healthy control (HC) groups. The distributions clearly illustrate a trend of decreasing frequency with increasing disease severity. Specifically, healthy controls exhibit the highest central frequency, indicating typical neural oscillatory activity. Subjects with MCI show slightly reduced frequency values, signifying the onset of neural slowing, while the AD group displays the lowest central frequencies, reflecting significant neural slowing commonly observed in Alzheimer's pathology. This consistent pattern across multiple critical EEG channels underscores the physiological relevance of BayesENDS' latent dynamics. The clear association between disease severity and decreased oscillatory frequency validates the biological interpretability of our model, highlighting its potential utility for understanding Alzheimer's disease progression and supporting clinical decision-making.

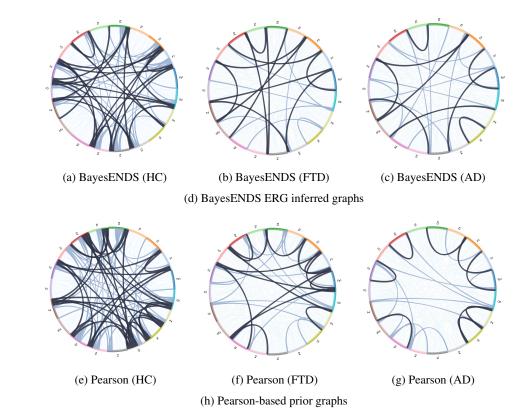


Figure 3: Comparison of EEG connectivity graphs inferred by BayesENDS versus Pearson correlation-based priors across healthy controls (HC), frontotemporal dementia (FTD), and Alzheimer's disease (AD) groups.

#### 6.3 GRAPH CONNECTIVITY VISUALIZATION

We also visualized BayesENDS' inferred event-relational graphs (ERGs) against Pearson correlation-based connectivity graphs (Figure 3). Chord diagrams revealed distinct connectivity patterns across healthy controls (HC), FTD, and AD groups. BayesENDS' ERGs captured biologically plausible trends: HC showed dense, robust connectivity, while FTD and AD exhibited progressively sparser and weaker links. These patterns aligned with Pearson-derived graphs, validating the ERG's effectiveness. The coherence underscores ERG's ability to reflect dementia-related network degradation, strengthening model interpretability and clinical relevance.

# 7 CONCLUSION

We presented **BayesENDS**, a Bayesian electrophysiological neural dynamical system that infers latent event dynamics and a conditional event–relational graph directly from multichannel EEG. By coupling an Event Posterior Differential Equation (EPDE) with a Mean–Evolving Lognormal Process (MELP) and an electrophysiology-informed dLIF prior, the model yields identifiable, physiology-aware latents and supports end-to-end prediction. Our theory provides a tractable IVP-based upper bound for the event–prior KL and establishes stability of the inferred graph to lag noise, while experiments on synthetic and AD EEG data demonstrate superior accuracy and interpretable biomarkers.

# 8 ETHICS STATEMENT

Compliant with ICLR ethics. No human/animal subjects or IRB approval required. Public datasets used under license; no re-identification attempted. For academic use; real-world applications need further validation. No harmful insights, conflicts, or sponsorship. Possible dataset biases noted. Methods documented for reproducibility.

# 9 REPRODUCIBILITY STATEMENT

We have taken deliberate steps to facilitate the reproducibility of our results. A complete description of the model architectures, training procedures, and evaluation protocols is provided in the main text. All datasets used are publicly available and cited with their licenses.

#### REFERENCES

- Madhav Acharya, Ravinesh C. Deo, Xiaohui Tao, Prabal Datta Barua, Aruna Devi, Anirudh Atmakuru, and Ru-San Tan. Deep learning techniques for automated Alzheimer's and mild cognitive impairment disease using EEG signals: A comprehensive review of the last decade (2013–2024). *Computer Methods and Programs in Biomedicine*, 259:108506, 2025. doi: 10.1016/j.cmpb.2024.108506.
- Hamdi Altaheri, Ghulam Muhammad, and Mansour Alsulaiman. Physics-informed attention temporal convolutional network for eeg-based motor imagery classification. *IEEE transactions on industrial informatics*, 19(2):2249–2258, 2022.
- Claudio Babiloni et al. International federation of clinical neurophysiology (ifcn) recommendations on electroencephalography in Alzheimer's disease and other dementias. *Clinical Neurophysiology*, 2021.
- André M. Bastos and Jan-Mathijs Schoffelen. A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Frontiers in Systems Neuroscience*, 9:175, 2016. doi: 10.3389/fnsys.2015.00175.
- Guoqiang Q. Bi and Mu-Ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of Neuroscience*, 18(24):10464–10472, 1998. doi: 10.1523/JNEUROSCI.18-24-10464.1998.
- Anthony N. Burkitt. A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. *Biological Cybernetics*, 95(1):1–19, 2006. doi: 10.1007/s00422-006-0068-6.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Justin Dauwels, François Vialatte, and Andrzej Cichocki. A comparative study of synchrony measures for the early diagnosis of Alzheimer's disease based on eeg. *Clinical Neurophysiology*, 121 (1):81–94, 2010.
- Mohsen Ehteshamzad et al. Assessing the potential of EEG in early detection of Alzheimer's disease. *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring*, 2024. doi: 10.3233/ADR-230159.
- Daniel E. Feldman. The spike-timing dependence of plasticity. *Neuron*, 75(4):556–571, 2012. doi: 10.1016/j.neuron.2012.08.001.
- Mansura Habiba and Barak A Pearlmutter. Neural ordinary differential equation based recurrent neural network model. In 2020 31st Irish signals and systems conference (ISSC), pp. 1–6. IEEE, 2020.
- Hengguan Huang, Hongfu Liu, Hao Wang, Chang Xiao, and Ye Wang. Strode: Stochastic boundary ordinary differential equation. In *International Conference on Machine Learning*, pp. 4435–4445. PMLR, 2021.

- Cosimo Ieracitano, Nadia Mammone, Amir Hussain, and Francesco Carlo Morabito. A novel multi-modal machine learning based approach for automatic classification of eeg recordings in dementia. *Neural Networks*, 123:176–190, 2020. doi: 10.1016/j.neunet.2019.12.006.
  - Jaeseung Jeong. Eeg dynamics in patients with alzheimer's disease. *Clinical Neurophysiology*, 115 (7):1490–1505, 2004. doi: 10.1016/j.clinph.2004.01.001.
  - Pramod Kachare, Digambar Puri, Sandeep B Sangle, Ibrahim Al-Shourbaji, Abdoh Jabbari, Raimund Kirner, Abdalla Alameen, Hazem Migdady, and Laith Abualigah. Lcadnet: a novel light cnn architecture for eeg-based alzheimer disease detection. *Physical and Engineering Sciences in Medicine*, 47(3):1037–1050, 2024.
  - Seul-Kee Kim, Hayom Kim, Sang Hee Kim, Jung Bin Kim, and Laehyun Kim. Electroencephalography-based classification of alzheimer's disease spectrum during computer-based cognitive testing. *Scientific Reports*, 14, 2024. doi: 10.1038/s41598-024-55656-8.
  - Dominik Klepl, Min Wu, and Fei He. Graph neural network-based EEG classification: A survey. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 32:493–503, 2024. doi: 10.1109/TNSRE.2024.3355750.
  - Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and Brent J Lance. Eegnet: a compact convolutional neural network for eeg-based brain–computer interfaces. *Journal of neural engineering*, 15(5):056013, 2018.
  - Keyvan Mahjoory, Vadim V. Nikulin, Léon Botrel, Klaus Linkenkaer-Hansen, Mario Massimo Fato, and Stefan Haufe. Consistency of EEG source localization and connectivity estimates. *NeuroImage*, 152:590–601, 2017. doi: 10.1016/j.neuroimage.2017.02.076.
  - Christoph M. Michel and Denis Brunet. EEG source imaging: A practical review of the analysis steps. *Frontiers in Neurology*, 10:325, 2019. doi: 10.3389/fneur.2019.00325.
  - Christoph M. Michel, Micah M. Murray, G. Lantz, Sara Gonzalez Andino, Luigi Spinelli, and R. Grave de Peralta Menendez. EEG source imaging. *Clinical Neurophysiology*, 115(10):2195–2222, 2004. doi: 10.1016/j.clinph.2004.06.001.
  - Andreas Miltiadous, Katerina D. Tzimourta, Theodora Afrantou, Panagiotis Ioannidis, Nikolaos Grigoriadis, Dimitrios G. Tsalikakis, Pantelis Angelidis, Markos G. Tsipouras, Evripidis Glavas, Nikolaos Giannakeas, and Alexandros T. Tzallas. "a dataset of eeg recordings from: Alzheimer's disease, frontotemporal dementia and healthy subjects", 2024.
  - Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):61–74, 2019.
  - Aruane Mello Pineda, Fernando M. Ramos, Luiz Eduardo Betting, and Andriana S. L. O. Campanharo. Use of complex networks for the automatic detection and the diagnosis of Alzheimer's disease. In Ignacio Rojas, Gonzalo Joya, and Andreu Catala (eds.), *Advances in Computational Intelligence: 15th International Work-Conference on Artificial Neural Networks (IWANN 2019), Proceedings, Part I*, volume 11506 of *Lecture Notes in Computer Science*, pp. 115–126, Cham, 2019. Springer. doi: 10.1007/978-3-030-20521-8\_10.
  - Seyed-Ali Sadegh-Zadeh, Elham Fakhri, Mahboobe Bahrami, Elnaz Bagheri, Razieh Khamsehashari, Maryam Noroozian, and Amir M Hajiyavand. An approach toward artificial intelligence alzheimer's disease diagnosis using brain signals. *diagnostics*, 13(3):477, 2023.
  - Md Nurul Ahad Tawhid, Siuly Siuly, Enamul Kabir, and Yan Li. Advancing alzheimer's disease detection: a novel convolutional neural network based framework leveraging eeg data and segment length analysis. *Brain Informatics*, 12(1):13, 2025. doi: 10. 1186/s40708-025-00260-3. URL https://braininformatics.springeropen.com/articles/10.1186/s40708-025-00260-3.
  - Mário L. Vicchietti, Fernando M. Ramos, Luiz E. Betting, and Andriana S. L. O. Campanharo. Computational methods of eeg signals analysis for Alzheimer's disease classification. *Scientific Reports*, 13:8184, 2023. doi: 10.1038/s41598-023-32664-8.

Yihe Wang, Nadia Mammone, Darina Petrovsky, Alexandros T. Tzallas, Francesco C. Morabito, and Xiang Zhang. ADformer: A multi-granularity transformer for EEG-based Alzheimer's disease assessment. arXiv preprint arXiv:2409.00032, 2024.

Yihe Wang, Nan Huang, Nadia Mammone, Marco Cecchi, and Xiang Zhang. Lead: Large foundation model for eeg-based alzheimer's disease detection. *arXiv preprint arXiv:2502.01678*, 2025.

Xiang Zhang and Lina Yao. Deep learning for EEG-based brain-computer interfaces: representations, algorithms and applications. World Scientific, 2021.

# A LLM USAGE STATEMENT

In accordance with the ICLR 2026 submission guidelines, we disclose that large language models (LLMs) were used only for language editing and proofreading of this manuscript. No LLMs were employed for generating research ideas, designing methodologies, producing experimental results, or creating data. All scientific content, analysis, and conclusions were developed and verified by the authors. The authors take full responsibility for the integrity, originality, and accuracy of the submission.

# B PROOF OF THEOREM 2

Proof. By definition,

$$KL(q||p_r) = \int_0^S q(t) \log \frac{q(t)}{r(t)e^{-\int_0^t r}} dt.$$

Using  $m=-e^{-t}$ , we have  $t=-\log(-m)=M$  and  $dt=-\frac{1}{m}\,dm$ . Hence

$$KL(q||p_r) = \int_{-e^{-S}}^{0} \frac{q(M)}{M} \log \frac{q(M)}{Mr(M)e^{-\int_0^M r}} \left(-\frac{1}{m}\right) dm = \int_{-e^{-S}}^{0} g(m) dm.$$

This improper integral equals  $\lim_{\varepsilon \downarrow 0} \int_{-e^{-S}}^{-\varepsilon} g(m) \, dm = \lim_{\varepsilon \downarrow 0} G(-\varepsilon)$ , proving Eq. (17). Since  $r(t) \in [a,b]$  and q is integrable, g is locally integrable near  $m=0^-$ . Split the integral at  $-\varepsilon$  and at  $-2\varepsilon$ :

$$\int_{-e^{-S}}^{0} g = \int_{-e^{-S}}^{-\varepsilon} g + \int_{-\varepsilon}^{0} g \le G(-\varepsilon) + \left| \int_{-\varepsilon}^{0} g \right| \le G(-\varepsilon) + \left| \int_{-2\varepsilon}^{-\varepsilon} g \right| = G(-\varepsilon) + |G(-2\varepsilon) - G(-\varepsilon)|,$$

which yields Eq. (18). The last inequality uses Lemma 1 (applied to the IVPs  $G_1'(m) = g(m)$  and  $G_2'(m) = g(m)$  shifted by  $\varepsilon$ ) to control the tail near the singular endpoint and the fact that r is bounded in [a,b], ensuring g remains controlled as  $m \to 0^-$ . As  $\varepsilon \downarrow 0$ , the tail vanishes by dominated convergence, hence  $\mathcal{U}_{\varepsilon} \to \mathrm{KL}(q||p_r)$ .

## C PROOF OF THEOREM 3

*Proof.* (Lipschitzness). The absolute value is 1–Lipschitz:  $||x|-|y|| \leq |x-y|$ . The function  $u\mapsto e^{-\alpha u}$  on  $u\geq 0$  has derivative  $|-\alpha e^{-\alpha u}|\leq \alpha$ , hence it is  $\alpha$ –Lipschitz on  $\mathbb{R}_{\geq 0}$ . By composition of Lipschitz maps,

$$|\phi_{\alpha}(x) - \phi_{\alpha}(y)| = \left| e^{-\alpha|x|} - e^{-\alpha|y|} \right| \le \alpha \left| |x| - |y| \right| \le \alpha |x - y|,$$

establishing Eq. (19). Taking  $y = x + \xi$  gives Eq. (20).

(Averaging). Using linearity of the average and triangle inequality,

$$\left| \tilde{\bar{A}}_{ij} - \bar{A}_{ij} \right| = \left| \frac{1}{MS} \sum_{m,s} \left( \tilde{e}_{ij}(t_m; T^{(s)}) - e_{ij}(t_m; T^{(s)}) \right) \right| \le \frac{1}{MS} \sum_{m,s} \left| \tilde{e}_{ij} - e_{ij} \right| \le \frac{\alpha}{MS} \sum_{m,s} |\xi_{ij}(t_m; T^{(s)})|,$$

which is Eq. (21).

(Matrix bound). Define  $\Delta^{(m,s)} := [\tilde{e}_{ij}(t_m;T^{(s)}) - e_{ij}(t_m;T^{(s)})]_{i \neq j}$ , so  $\tilde{\bar{A}} - \bar{A} = \frac{1}{MS} \sum_{m,s} \Delta^{(m,s)}$ . By the triangle inequality of the Frobenius norm,

$$\|\widetilde{\bar{A}} - \bar{A}\|_{\mathrm{F}} \le \frac{1}{MS} \sum_{m,s} \|\Delta^{(m,s)}\|_{\mathrm{F}}.$$

Entry-wise inequality Eq. 20 implies  $\|\Delta^{(m,s)}\|_{\mathrm{F}} \leq \alpha \|\Xi^{(m,s)}\|_{\mathrm{F}}$ , giving Eq. (22).

# D PROOF OF COROLLARY 4

*Proof.* (i) From Eq. (21),  $\overline{|\xi_{ij}|} \leq \varepsilon_{\infty}$ , giving the entry-wise claim. For the matrix bound,  $\|\Xi^{(m,s)}\|_{\rm F} \leq \varepsilon_{\infty} \sqrt{C(C-1)}$  for every (m,s); apply Eq. (22).

(ii) For each fixed (i,j), define i.i.d. variables  $Y_{m,s} := \tilde{e}_{ij}(t_m;T^{(s)}) - e_{ij}(t_m;T^{(s)})$ . By Eq. (20),  $Y_{m,s}$  is an  $\alpha$ -Lipschitz transform of  $\xi_{ij}(t_m;T^{(s)})$ , hence  $Y_{m,s}$  is sub-Gaussian with proxy  $\alpha^2\sigma^2$  (standard Lipschitz preservation of the  $\psi_2$ -norm). Then  $\Delta_{ij} = (MS)^{-1} \sum_{m,s} Y_{m,s}$  is sub-Gaussian with proxy  $\alpha^2\sigma^2/(MS)$ , yielding Eq. (24) via the Chernoff bound.

For the Gaussian–mean bound, use  $|Y_{m,s}| \leq \alpha |\xi_{ij}(t_m;T^{(s)})|$  and linearity:

$$\mathbb{E}|\Delta_{ij}| \le \frac{\alpha}{MS} \sum_{m,s} \mathbb{E}|\xi| = \alpha \,\mathbb{E}|\xi| = \alpha \,\sigma \sqrt{\frac{2}{\pi}},$$

where the last equality is the mean absolute value of a zero–mean Gaussian. Summing the entry–wise bounds in quadrature gives the Frobenius expectation in Eq. (25).

# E ALZHEIMER'S DISEASE EEG BASELINE METHODS

To rigorously evaluate our proposed method, we benchmarked it against several representative deep learning approaches commonly utilized for EEG analysis. These baselines include convolutional, recurrent, attention-based, and transformer-based models, each demonstrating distinct strengths for capturing various aspects of EEG signal patterns.

**EEGNet** (Lawhern et al., 2018) is a compact convolutional neural network initially developed for EEG-based brain–computer interfaces. It integrates depthwise and separable convolutions to effectively capture temporal, spatial, and frequency-specific characteristics in EEG data, making it a well-recognized lightweight yet powerful model in EEG classification tasks.

**LCADNet** (Kachare et al., 2024) is specifically tailored for Alzheimer's disease detection from EEG data. Utilizing optimized convolutional structures designed for computational efficiency without sacrificing discriminative power, LCADNet achieves competitive performance in resource-limited environments, making it a strong baseline for EEG-based AD diagnosis.

**LSTM** (Zhang & Yao, 2021) embodies recurrent neural networks tailored for modeling temporal dependencies inherent in EEG signals. By maintaining and updating hidden states across sequences, LSTMs effectively capture long-term dynamics and temporal correlations, making them naturally suitable for sequential EEG analyses.

**ATCNet** (Altaheri et al., 2022) employs a physics-informed architecture combining temporal convolutions with attention mechanisms. Originally proposed for motor imagery EEG classification, it effectively captures both local temporal details and global dependencies, showcasing adaptability across various EEG applications.

**ADformer** (Wang et al., 2024) is a multi-granularity transformer specifically crafted for Alzheimer's disease evaluation using EEG signals. It utilizes multi-scale attention mechanisms to concurrently model fine-grained and coarse-grained temporal information, setting a high-performance standard in EEG-based AD diagnostics.

**LEAD** (Wang et al., 2025) exemplifies the recent advancement toward large-scale foundation models in EEG analysis. Pre-trained extensively on vast EEG datasets and fine-tuned for Alzheimer's disease detection, LEAD leverages transfer learning to provide robust, generalizable EEG representations, establishing a new benchmark in EEG-based clinical assessments.

#### E.1 DATASETS

## E.1.1 TOY DATASET AND DATA GENERATION

Frequency bands and sampling. To systematically evaluate our model's ability to capture latent event dynamics, we constructed synthetic datasets with clearly defined frequency bands. We generated latent event rates  $\lambda$  from truncated normal distributions centered at the midpoint of each target frequency band: low band [5–10 Hz] with  $\lambda \sim \text{TruncNormal}(\mu = 7.5, \sigma = 1.0; [5, 10])$ , middle band [10–15 Hz] with  $\lambda \sim \text{TruncNormal}(\mu = 12.5, \sigma = 1.0; [10, 15])$ , and high band [15–20 Hz] with  $\lambda \sim \text{TruncNormal}(\mu = 17.5, \sigma = 1.0; [15, 20])$ . This design ensures concentrated event-rate distributions within each band while avoiding frequencies outside the desired range.

**Data scale and splitting strategy.** For each frequency band, we independently generated three data splits: a training set with 150 distinct event rates, each having 50 sequences (7,500 sequences total); a validation set with 25 new event rates and 50 sequences per rate (1,250 sequences total); and a test set with an additional 25 new event rates, again with 50 sequences per rate (1,250 sequences total). Importantly, no overlap of event rates occurs across training, validation, and test splits to ensure proper evaluation of model generalization.

**Sequence generation.** Each synthetic sequence comprises 20 observations, constructed by sampling inter-event times  $\Delta t_i$  from an exponential distribution with parameter  $\lambda$ . The event timestamps  $t_i$  are obtained cumulatively by  $t_i = \sum_{j=1}^i \Delta t_j$ . Observations  $y_i$  are subsequently generated using the relationship:

$$y_i = \sin(t_i) + \eta_i, \quad \eta_i \sim \mathcal{N}(0, \sigma_\eta^2),$$

where the default noise level is  $\sigma_{\eta}=0.07$ . Additional sensitivity analyses varied  $\sigma_{\eta}$  within  $\{0.05, 0.10, 0.15\}$  to assess model robustness.

**Evaluation methodology.** Model performance was comprehensively evaluated using three criteria: (1) classification score (CS) assessing sequence-level predictive accuracy, (2) uncertainty calibration, quantified through the median and 95% confidence interval of the estimated event rate  $\hat{\lambda}$ , obtained via nonparametric resampling within each frequency band, and (3) structural fidelity, measured using intersection-over-union (IoU) between the predicted latent structure and the ground-truth event patterns.

#### E.1.2 ALZHEIMER'S DISEASE EEG DATASET

**Dataset AD Cohort A** (Miltiadous et al., 2024) consists of resting-state, eyes-closed EEG recordings from a total of 88 participants, categorized into 36 individuals diagnosed with Alzheimer's disease (AD), 23 patients with frontotemporal dementia (FTD), and 29 healthy control subjects (HC). The EEG data were collected using 19 electrodes arranged according to the international 10–20 placement system. The recordings have a sampling rate of 500 Hz and an average duration ranging from approximately 12 to 14 minutes per subject. Provided in adherence to the Brain Imaging Data Structure (BIDS) standard, the dataset includes both raw and preprocessed EEG signals, enabling robust comparative analysis across different dementia subtypes.

**Dataset AD Cohort B** (Sadegh-Zadeh et al., 2023) includes resting-state EEG data from 168 participants, segmented into 59 moderate Alzheimer's disease patients (AD), 7 individuals diagnosed with mild cognitive impairment (MCI), and 102 healthy controls (HC). EEG recordings were acquired using the standardized 10–20 electrode placement system, with data presented in MATLAB (.mat) format. Accompanying the EEG data are Mini-Mental State Examination (MMSE) scores, providing cognitive assessments for participants. This dataset is particularly tailored for the distinction of AD from MCI, thus serving as a valuable resource for investigations aimed at early Alzheimer's disease diagnosis.

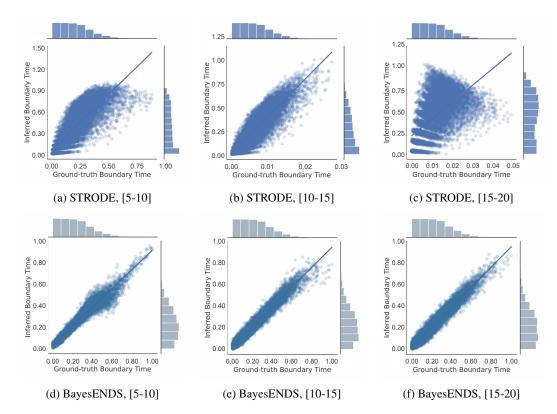


Figure 4: Predicted vs. ground-truth boundary times across frequency bands ([5–10], [10–15], [15–20] Hz): STRODE vs. BayesENDS.

#### E.2 TOY DATASET EXPERIMENTS

The toy dataset experiments evaluated how well BayesENDS and baseline neural ODE methods (NODE (Chen et al., 2018), ODE-RNN (Habiba & Pearlmutter, 2020), and STROD (Huang et al., 2021)) recover latent event dynamics across distinct frequency bands ([5–10], [10–15], and [15–20] Hz). As summarized in Table 3, baseline methods achieved strong cosine similarity (CS) scores across all frequency bands, reflecting good sequence-level prediction performance. However, their intersection-over-union (IoU) scores were uniformly zero, indicating a fundamental limitation in capturing the latent event structure. This outcome highlights the common issue with purely data-driven neural approaches: despite excellent predictive accuracy, they often fail to recover the underlying generative mechanisms of data.

In contrast, BayesENDS maintained similarly high CS scores while notably achieving meaning-ful IoU values (e.g., 0.473 at [5–10] Hz, 0.289 at [10–15] Hz, and 0.202 at [15–20] Hz). These non-zero IoU scores demonstrate that BayesENDS successfully captures latent structures consistent with the generative process, particularly emphasizing the method's ability to infer interpretable and biologically plausible event dynamics. Moreover, the predicted event rates from BayesENDS showed wider but informative uncertainty intervals, aligning closely with the true frequency band intervals. Such uncertainty quantification highlights BayesENDS' capacity to provide both accurate and interpretable latent dynamics in noisy and ambiguous settings.

#### E.3 FIGURES FOR ALBATION EXPERIMENTS

Table 3: Toy dataset results by frequency band. CS: Cosine Similarity, IoU: Intersection-over-Union.

Frequency Band (Hz)	Model	CS	Median Rate	95% CI	IoU
[5, 10]	NODE	0.951	1.000	[1.000, 1.000]	0.000
[5, 10]	ODE-RNN	0.951	1.000	[1.000, 1.000]	0.000
[5, 10]	STRODE	0.967	0.340	[0.269, 0.410]	0.000
[5, 10]	BayesENDS (Ours)	0.983	7.532	[4.300, 14.867]	0.473
[10, 15]	NODE	0.951	1.000	[1.000, 1.000]	0.000
[10, 15]	ODE-RNN	0.951	1.000	[1.000, 1.000]	0.000
[10, 15]	STRODE	0.964	0.251	[0.153, 0.348]	0.000
[10, 15]	BayesENDS (Ours)	0.982	12.503	[7.587, 24.918]	0.289
[15, 20]	NODE	0.951	1.000	[1.000, 1.000]	0.000
[15, 20]	ODE-RNN	0.951	1.000	[1.000, 1.000]	0.000
[15, 20]	STRODE	0.961	0.532	[0.369, 0.695]	0.000
[15, 20]	BayesENDS (Ours)	0.976	18.843	[10.465, 35.244]	0.202