

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 BAYEENDS: 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., 2024b; 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 neuron 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 infers event-driven latent dynamics and a conditional interaction graph directly from multichannel EEG. Concretely, 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

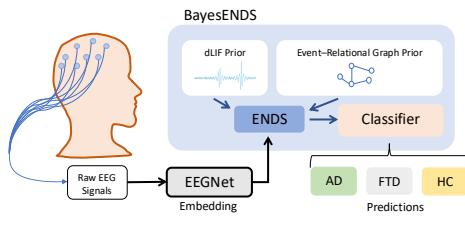


Figure 1: Overview of the BayesENDS pipeline

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 under a variational inference framework; 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). Recent large-scale and representation-learning approaches for AD-EEG continue this predominantly discriminative perspective: LEAD builds a foundation model for AD detection from multi-dataset EEG (Wang et al., 2025), COMET introduces hierarchical contrastive learning for medical time series including AD-EEG (Wang et al., 2023), Medformer employs multi-granularity Transformer patching for disease classification with AD cohorts (Wang et al., 2024a), and manifold-based vector-field modeling reconstructs high-density AD-EEG dynamics from low-density recordings (Peach et al., 2023). 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 with sequence-level labels. Given a dataset $\mathcal{D} = \{(X^{(n)}, Y^{(n)})\}_{n=1}^N$ 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 $c \in \{1, \dots, C\}$ and sequence n , let $T_c^{(n)} = \{t_{c,k}^{(n)}\}_{k=1}^{K_c^{(n)}}$ denote the (unknown) latent event times, where $t_{c,k}^{(n)}$ is the k -th event time on channel c in sequence n , and $K_c^{(n)}$ is the (latent) number of events on that channel. 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 the conditional graph dynamics $P(G^{(n)}(\cdot) \mid \mathbf{p}^{(n)}(\cdot))$ 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.

108 The interaction strength between channels is modeled as a function of the temporal *co-occurrence*
 109 and *ordering* of inferred events. This assumption is grounded in neurobiological mechanisms of
 110 *spike-timing-dependent plasticity* (STDP), where near-coincident pre- and post-synaptic spikes mod-
 111 ulate synaptic efficacy, forming the basis for capturing dynamic interactions between different neural
 112 regions(Bi & Poo, 1998; Feldman, 2012).

4 BAYESIAN NEURAL DYNAMICAL SYSTEM

4.1 OVERVIEW

118 We introduce BayesENDS, a Bayesian neural dynamical system for multichannel sequences that
 119 represents each channel with a latent *event* process and couples channels through a conditional
 120 *event-relational graph* (ERG) driven by the timing and ordering of inferred events. This design
 121 is motivated by settings where the clinically relevant signal resides in *when* events occur and *how*
 122 they align across channels. In EEG for Alzheimer’s disease, for instance, oscillatory slowing and
 123 disrupted coordination suggest that event timing and cross-channel alignment are predictive, while
 124 neurobiological plasticity links near-coincident spikes to stronger coupling.

125 At a high level, BayesENDS consists of three interacting components. First, an *event posterior*
 126 *differential equation* (EPDE) summarizes per-channel event dynamics by producing posterior distri-
 127 butions over latent event times $\{T_c^{(n)}\}$ given the observed multichannel sequence $X^{(n)}$. Second, a
 128 *mean-evolving lognormal mechanism* (MELP) uses EPDE outputs as mean parameters to generate
 129 stochastic inter-event timing between successive latent event times $t_{c,k}^{(n)}$, ensuring positive and flex-
 130 ible (potentially multimodal) timing statistics. Third, an *event-relational graph* $G^{(n)}(t)$ is inferred
 131 from event co-occurrence and cross-channel lags derived from $T^{(n)}$, via an STDP-shaped mapping
 132 that encodes how the timing of events on one channel modulates effective coupling to others.

134 For each labeled sequence $(X^{(n)}, Y^{(n)})$, the triple $(X^{(n)}, T^{(n)}, G^{(n)})$ is passed to a decoder
 135 $p_\theta(Y^{(n)} \mid X^{(n)}, T^{(n)}, G^{(n)})$ for downstream prediction (e.g., AD vs. control). Training is end-
 136 to-end via variational learning that jointly optimizes the EPDE and MELP while learning ERG
 137 dynamics under weak regularization; further details are given in the learning subsection.

4.2 LEARNING

141 We train BayesENDS end-to-end through variational inference by *minimizing* the negative evidence
 142 lower bound (ELBO). For labeled data $\{(X^{(n)}, Y^{(n)})\}_{n=1}^N$, let $T^{(n)}$ collect all latent event times
 143 $\{T_c^{(n)}\}_{c=1}^C$ and let $\tau^{(n)}$ collect the corresponding inter-event intervals $\{\tau_{c,k}^{(n)}\}$. The EPDE induces
 144 an approximate posterior $q_\phi(T^{(n)} \mid X^{(n)})$, the MELP defines $q_\phi(\tau^{(n)} \mid X^{(n)})$, and the decoder
 145 is $p_\theta(Y^{(n)} \mid X^{(n)}, T^{(n)}, G^{(n)})$, where $G^{(n)}$ is the ERG associated with $X^{(n)}$ and η denotes ERG
 146 parameters.

147 We write the ELBO as

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

$$- \text{KL}_T^{(n)} - \text{KL}_\tau^{(n)} \quad (2)$$

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

155 where $\text{KL}_T^{(n)} := \text{KL}(q_\phi(T^{(n)} \mid X^{(n)}) \parallel p_{\text{dLIF}}(T))$ compares the EPDE-induced path law to the
 156 electrophysiology-informed event prior $p_{\text{dLIF}}(T)$, and $\text{KL}_\tau^{(n)} := \text{KL}(q_\phi(\tau^{(n)} \mid X^{(n)}) \parallel p_0(\tau))$ pe-
 157 neralizes deviation from a lognormal(-mixture) prior over inter-event intervals. The term $\mathcal{R}_{\text{LIF}}^{(n)}$ softly
 158 enforces leaky-integrate-and-fire consistency on differentiable rate proxies read out from the EPDE
 159 state, with weight $\lambda_{\text{LIF}} \geq 0$. The term $\mathcal{R}_{\text{ERG}}^{(n)}$ is a weak, observable-based regularizer that nudges
 160 ERG edges toward experimental statistics computed from $X^{(n)}$ (e.g., correlation-based summaries),
 161 with strength $\beta \geq 0$.

Challenges. Three technical issues arise in optimizing equation 3. First, $\text{KL}_T^{(n)}$ involves *path measures* induced by a differential equation and is intractable in closed form (it integrates over an infinite-dimensional trajectory); we therefore replace it with a tractable integral–rate surrogate that depends only on the dLIF rate $r(t)$, with a formal bound given in the Theory subsection. Second, enforcing a LIF prior directly is difficult because the spike function in LIF is *non-differentiable*, which prevents straightforward use in gradient-based training; instead we introduce differentiable rate proxies and constrain them to follow dLIF laws via $\mathcal{R}_{\text{LIF}}^{(n)}$. Third, ERG learning lacks ground-truth edges; to avoid over-constraining the graph, we only use the weak regularizer $\mathcal{R}_{\text{ERG}}^{(n)}$ to bias edge strengths toward experimental observables, leaving the fine-grained graph structure to be driven by event lags inferred from the EPDE–MELP posterior.

4.3 PRIOR: ELECTROPHYSIOLOGY–INFORMED DLIF PRIOR

We place a biophysical prior on latent event timing by instantiating each channel’s latent events $T_c^{(n)} = \{t_{c,k}^{(n)}\}$ as a renewal process whose hazard is derived from a differentiable leaky–integrate–and–fire (dLIF) abstraction (Burkitt, 2006). For channel c , the (rescaled) membrane potential evolves as

$$\frac{d}{dt}u_c(t) = b_c(t) - u_c(t), \quad b_c(t) > 1, \quad (4)$$

where $b_c(t)$ is an effective (learned) input drive. Given this membrane dynamics, the implied instantaneous *firing rate* is

$$r_c(t) = \left[-\log(1 - 1/b_c(t)) \right]^{-1}. \quad (5)$$

This rate induces a dLIF inter-event time density

$$p_{\text{dLIF},c}(t) = r_c(t) \exp\left(-\int_0^t r_c(s) ds\right), \quad (6)$$

and the resulting dLIF prior for channel c is the renewal law $p_{\text{dLIF}}(T_c^{(n)})$ with hazard $r_c(t)$. We parameterize $b_c(t)$ by a bounded neural mapping from learned embeddings, for example

$$b_c(t) = 1 + \text{softplus}(g_\xi(z_c^{(n)}(t))), \quad (7)$$

where $z_c^{(n)}(t)$ denotes features derived from $X^{(n)}$. This construction guarantees $b_c(t) > 1$ and thus $r_c(t) > 0$. Absolute and refractory effects are incorporated through a smooth gating factor $\alpha_c^{(n)}(t) \in (0, 1]$ constructed from recent events in $T_c^{(n)}$, using $r_c(t) = \alpha_c^{(n)}(t) r_c(t)$ 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

$$\mathcal{R}_{\text{LIF}} = \sum_c \int_0^S (\hat{r}_c^{(n)}(t) - r_c(t))^2 dt, \quad (8)$$

where $\hat{r}_c^{(n)}(t)$ is a differentiable rate proxy read from the EPDE state for sequence n over a time horizon $[0, S]$. This encourages the learned rates to follow dLIF membrane dynamics without invoking non-differentiable spike functions (Neftci et al., 2019). The variational KL between the EPDE–induced path law $q_\phi(T_c^{(n)} \mid X^{(n)})$ and $p_{\text{dLIF}}(T_c^{(n)})$ 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 sequence $(X^{(n)}, Y^{(n)})$ and channel c , let $q_c^{(n)}(t \mid x_c^{(n)})$ denote the density of the next event time given the observed channel signal $x_c^{(n)}$. If $\tilde{t}_{c,k-1}^{(n)}$ denotes the previous (predicted) event time on that channel, the expected next event time is

$$\tilde{t}_{c,k}^{(n)} = \int_{\tilde{t}_{c,k-1}^{(n)}}^{\infty} t q_c^{(n)}(t \mid x_c^{(n)}) dt. \quad (9)$$

To express this update via an initial value problem (IVP), we introduce an auxiliary function $\Phi_c^{(n)}(t)$ whose derivative accumulates the contribution of $q_c^{(n)}$:

$$(\Phi_c^{(n)})'(t) = -t q_c^{(n)}(t | x_c^{(n)}). \quad (10)$$

Its initial value encodes the full expectation under $q_c^{(n)}$:

$$\Phi_c^{(n)}(0) = \int_0^\infty t q_c^{(n)}(t | x_c^{(n)}) dt. \quad (11)$$

With this definition, the expected next event time in equation 9 can be written as the IVP solution evaluated at the previous event time:

$$\tilde{t}_{c,k}^{(n)} = \Phi_c^{(n)}(\tilde{t}_{c,k-1}^{(n)}). \quad (12)$$

Directly solving equation 10–12 and computing $\Phi_c^{(n)}(0)$ is intractable, so we approximate this mapping with a differentiable neural surrogate that updates the predicted next event time:

$$\tilde{t}_{c,k}^{(n)} = f_{\theta_\Phi}(\tilde{t}_{c,k-1}^{(n)}, x_c^{(n)}), \quad (13)$$

implemented to ensure $\tilde{t}_{c,k}^{(n)} > \tilde{t}_{c,k-1}^{(n)}$, so that latent events remain strictly ordered in time.

Consequently, differentiating the ideal $\Phi_c^{(n)}(t)$ with respect to t yields an *event-time posterior* of the form

$$q_c^{(n)}(t | x_c^{(n)}) = -\frac{(\Phi_c^{(n)})'(t)}{t}, \quad (14)$$

which we approximate with the EPDE parameterization. Across channels, the family $\{q_c^{(n)}(t | x_c^{(n)})\}_{c=1}^C$ provides a parametric approximation to the event-time posteriors $\{p_c^{(n)}(t)\}_{c=1}^C$ introduced in the problem formulation, and jointly defines the EPDE-induced posterior $q_\phi(T^{(n)} | X^{(n)})$ over latent event times $T^{(n)} = \{T_c^{(n)}\}_{c=1}^C$ used in the variational objective. The predicted next-event times $\tilde{t}_{c,k}^{(n)}$ then serve as mean parameters for the mean–evolving lognormal mechanism (MELP) that models stochastic variability in event timing, as detailed in the next subsection.

4.5 SAMPLING: MEAN–EVOLVING LOGNORMAL PROCESS (MELP)

For each sequence $X^{(n)}$ and channel c , given the previous event time $\tilde{t}_{c,i-1}^{(n)}$, the EPDE produces a K -dimensional vector of *candidate mean intervals* $\tilde{\tau}_c^{(n)} = (\tilde{\tau}_{c,1}^{(n)}, \dots, \tilde{\tau}_{c,K}^{(n)}) \in \mathbb{R}_+^K$, together with mixture weights $\mathbf{w}_c^{(n)} = (w_{c,1}^{(n)}, \dots, w_{c,K}^{(n)}) \in \Delta^{K-1}$ and scales $\mathbf{s}_c^{(n)} = (s_{c,1}^{(n)}, \dots, s_{c,K}^{(n)}) \in \mathbb{R}_+^K$. MELP draws the inter–event interval $\tau_{c,i}^{(n)}$ from a lognormal mixture:

$$p(\tau_{c,i}^{(n)} | \tilde{t}_{c,i-1}^{(n)}, X^{(n)}) = \sum_{j=1}^K w_{c,j}^{(n)} \text{LogN}(\tau_{c,i}^{(n)}; \mu_{c,j}^{(n)}, (s_{c,j}^{(n)})^2), \quad (15)$$

where $\text{LogN}(\cdot; \mu, s^2)$ denotes a lognormal density with log-mean μ and log-variance s^2 .

For each component j , we choose $\mu_{c,j}^{(n)}$ as a function of $\tilde{\tau}_{c,j}^{(n)}$ and $s_{c,j}^{(n)}$ so that the *mean* of the corresponding lognormal distribution matches the EPDE-predicted interval $\tilde{\tau}_{c,j}^{(n)}$. This ties the mixture components’ average inter-event times directly to the EPDE outputs, while the scales $s_{c,j}^{(n)}$ control uncertainty around these means.

Sampling from MELP is reparameterized to keep gradients pathwise:

$$k \sim \text{Cat}(\mathbf{w}_c^{(n)}), \quad \varepsilon \sim \mathcal{N}(0, 1), \quad (16)$$

$$\tau_{c,i}^{(n)} = \exp(\mu_{c,k}^{(n)} + s_{c,k}^{(n)} \varepsilon), \quad t_{c,i}^{(n)} = t_{c,i-1}^{(n)} + \tau_{c,i}^{(n)}. \quad (17)$$

During training we use a differentiable variant of equation 16 (e.g., Gumbel–Softmax) and take hard samples at test time. MELP guarantees positive inter-event intervals, captures multimodal timing statistics, and yields closed-form component-wise KL terms against a lognormal(-mixture) prior in the learning objective. Moreover, the mixture expectation $\mathbb{E}[\tau_{c,i}^{(n)}] = \sum_{j=1}^K w_{c,j}^{(n)} \tilde{\tau}_{c,j}^{(n)}$ is available in closed form and is used in downstream computations such as computing ERG lags.

270 Table 1: Toy dataset results by frequency band. CS: Cosine Similarity, IoU: Intersection-over-Union.
271

272 Frequency Band (Hz)	273 Model	274 CS	275 Median Rate	276 95% CI	277 IoU
278 [5, 10]	279 NODE	280 0.951	281 1.000	282 [1.000, 1.000]	283 0.000
284 [5, 10]	285 ODE-RNN	286 0.951	287 1.000	288 [1.000, 1.000]	289 0.000
290 [5, 10]	291 STRODE	292 0.967	293 0.340	294 [0.269, 0.410]	295 0.000
296 [5, 10]	297 BayesENDS (Ours)	298 0.983	299 7.532	300 [4.300, 14.867]	301 0.473
302 [10, 15]	303 NODE	304 0.951	305 1.000	306 [1.000, 1.000]	307 0.000
308 [10, 15]	309 ODE-RNN	310 0.951	311 1.000	312 [1.000, 1.000]	313 0.000
314 [10, 15]	315 STRODE	316 0.964	317 0.251	318 [0.153, 0.348]	319 0.000
320 [10, 15]	321 BayesENDS (Ours)	322 0.982	323 12.503	324 [7.587, 24.918]	325 0.289
326 [15, 20]	327 NODE	328 0.951	329 1.000	330 [1.000, 1.000]	331 0.000
332 [15, 20]	333 ODE-RNN	334 0.951	335 1.000	336 [1.000, 1.000]	337 0.000
338 [15, 20]	339 STRODE	340 0.961	341 0.532	342 [0.369, 0.695]	343 0.000
344 [15, 20]	345 BayesENDS (Ours)	346 0.976	347 18.843	348 [10.465, 35.244]	349 0.202

285
286 4.6 EVENT-RELATIONAL GRAPH (ERG)
287288 We weakly bias $\bar{A}^{(n)}$ toward observable statistics computed from $X^{(n)}$ (e.g., Pearson correlations
289 $s_{ij}^{(n)}$ between channels i and j) via a simple Fisher- z alignment. We map both the observed and
290 ERG-implied correlations into z -space:

291
$$z_{ij}^{\text{obs},(n)} = \text{atanh}(s_{ij}^{(n)}), \quad z_{ij}^{\text{pred},(n)} = \text{atanh}(2\bar{A}_{ij}^{(n)} - 1), \quad (18)$$

292 and define the ERG regularizer as
293

294
$$\mathcal{R}_{\text{ERG}}^{(n)} = \sum_{i < j} \left[\frac{(z_{ij}^{\text{obs},(n)} - z_{ij}^{\text{pred},(n)})^2}{2\sigma^2} + \frac{1}{2} \log \sigma^2 \right], \quad (19)$$

295 where $\sigma > 0$ is a (fixed or globally learned) scale parameter controlling the strength of the alignment.
296 This regularizer (weighted by β in the learning objective) encourages consistency between ERG-
297 implied connectivity and experimental statistics, while still allowing the detailed edge structure to
298 be driven primarily by event lags inferred from the EPDE-MELP posterior. The implementation
299 details and additional theoretical results are provided in Appendix G and B, respectively.
300302 5 EXPERIMENTS
303304 In our experimental evaluations, we rigorously assessed the performance of BayesENDS across syn-
305 synthetic benchmarks and real Alzheimer’s disease (AD) EEG datasets. First, we validated BayesENDS
306 using synthetic event-sequence data, demonstrating its effectiveness in accurately inferring latent
307 event dynamics compared to baseline neural ODE models. Subsequently, we conducted compre-
308 hensive experiments on two diverse EEG datasets (AD cohorts A and B), covering Alzheimer’s
309 disease, frontotemporal dementia, mild cognitive impairment, and healthy controls.
310311 5.1 TOY DATASET
312313 The toy dataset experiments evaluated how well BayesENDS and baseline neural ODE methods
314 (NODE (Chen et al., 2018), ODE-RNN (Habiba & Pearlmuter, 2020), and STRODE (Huang et al.,
315 2021)) recover latent event dynamics across distinct frequency bands ([5–10], [10–15], and [15–20]
316 Hz). As summarized in Table 1, baseline methods achieved strong cosine similarity (CS) scores
317 across all frequency bands, reflecting good sequence-level prediction performance. However, their
318 intersection-over-union (IoU) scores were uniformly zero, indicating a fundamental limitation in
319 capturing the latent event structure. This outcome highlights the common issue with purely data-
320 driven neural approaches: despite excellent predictive accuracy, they often fail to recover the under-
321 lying generative mechanisms of data.
322323 In contrast, BayesENDS maintained similarly high CS scores while notably achieving meaningful
324 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-
325 zero IoU scores demonstrate that BayesENDS successfully captures latent structures consistent with

324
325
326 Table 2: Results on Alzheimer’s EEG datasets.
327
328
329
330
331
332
333
334

Model	Dataset AD cohort A		Dataset AD cohort B	
	Accuracy (%)	F1 (%)	Accuracy (%)	F1 (%)
EEGNet	68.10	66.49	71.37	60.85
LCADNet	70.52	68.12	72.44	49.38
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
LEAD	72.68	69.98	80.00	62.21
BayesENDS	75.03	72.69	89.82	64.87

335
336 Table 3: Ablation of spike-informed and connectome priors in **BayesENDS** on Alzheimer’s EEG
337 Datasets. Accuracy and F1 (%) reported as mean \pm s.d. across runs.
338
339
340
341
342
343
344
345
346
347
348

Dataset	Variant	Accuracy (%)	F1 (%)
AD cohort A	No prior	70.52 ± 11.83	65.46 ± 13.10
	dLIF prior	73.92 ± 9.84	71.41 ± 10.72
	ERG prior	72.75 ± 6.63	70.35 ± 7.91
	Dual priors	75.03 ± 8.29	72.69 ± 8.16
AD cohort B	No prior	83.22 ± 15.10	60.72 ± 12.37
	dLIF prior	87.98 ± 8.09	62.95 ± 9.24
	ERG prior	86.20 ± 9.96	65.63 ± 9.17
	Dual priors	89.82 ± 8.39	64.87 ± 10.67

349 the generative process, particularly emphasizing the method’s ability to infer interpretable and bio-
350 logically plausible event dynamics. Moreover, the predicted event rates from BayesENDS showed
351 wider but informative uncertainty intervals, aligning closely with the true frequency band intervals.
352 Such uncertainty quantification highlights BayesENDS’ capacity to provide both accurate and inter-
353 pretiable latent dynamics in noisy and ambiguous settings. Additional data generation protocol are
354 provided in Appendix F.

355
356

5.2 ALZHEIMER’S DISEASE EEG DATASET

357 We extensively evaluated BayesENDS on two diverse EEG datasets covering Alzheimer’s disease
358 (AD), frontotemporal dementia (FTD), mild cognitive impairment (MCI), and healthy controls, as
359 shown in Tabel 2. Across both datasets, BayesENDS outperformed state-of-the-art baselines such
360 as CNNs, RNNs, and transformers. In Dataset AD cohort A, BayesENDS showed a clear ability
361 to distinguish AD, FTD, and healthy participants, with consistently lower performance variabil-
362 ity across runs. This stability highlights its effectiveness in capturing subtle neural dynamics de-
363 spite EEG heterogeneity. In the more unbalanced Dataset AD cohort B, where distinctions between
364 moderate AD, MCI, and healthy controls are subtler, BayesENDS still achieved the highest diag-
365 nostic accuracy. While some baselines showed significant performance fluctuations, BayesENDS
366 remained robust, balancing sensitivity and specificity—demonstrating resilience to noise and distri-
367 bution shifts common in real-world EEG data. Overall, BayesENDS effectively extracts clinically
368 meaningful biomarkers from EEG, confirming its practical potential for accurate, interpretable AD
369 diagnostics. Additional dataset descriptions, baseline methods, and experimental setup are presented
370 in Appendix F.

371
372

5.3 ABLATION STUDIES

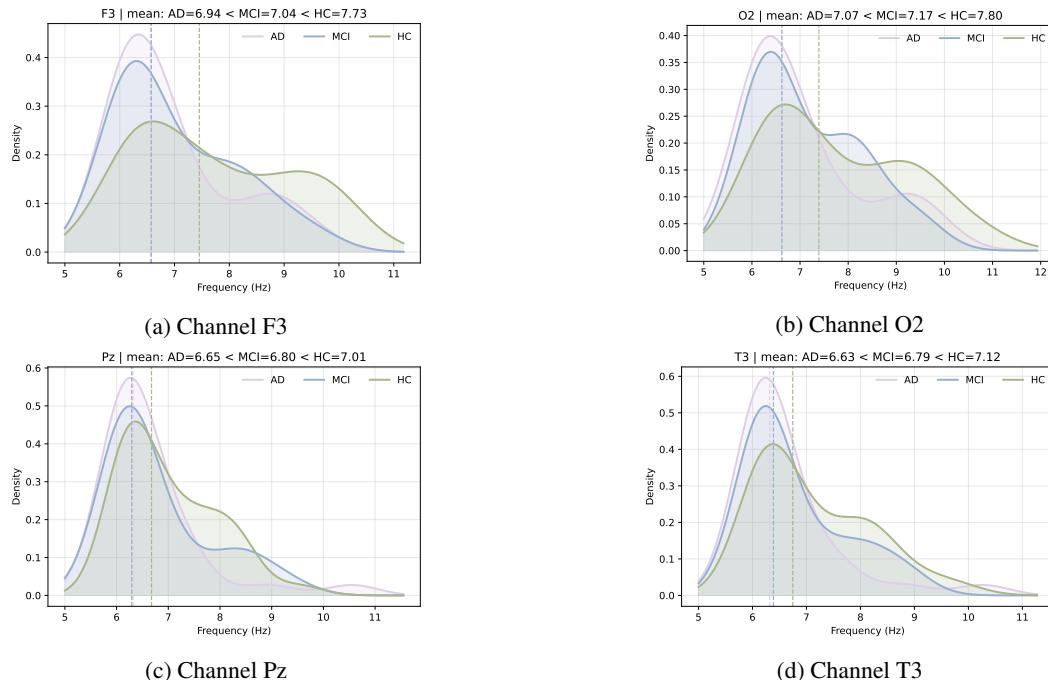
373 We ablated BayesENDS to assess each prior’s impact on accuracy and interpretability (Table 3).
374 On AD Cohort A, the no-prior baseline is moderate; adding the dLIF prior boosts accuracy, while
375 the ERG prior improves accuracy, and F1 via inter-channel modeling. Combining both yields the
376 best scores. On AD Cohort B, the pattern holds with larger gains: the dLIF prior gives the biggest
377 accuracy lift under subtler classes, and the ERG prior raises F1. Their combination again performs
best. Overall, each prior helps, and together they provide robust predictions on realistic EEG.

378 6 VISUALIZATIONS

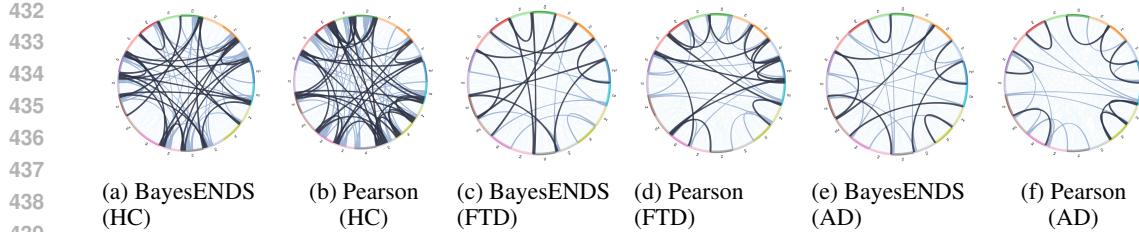
380
 381 Our visualizations emphasized BayesENDS' explainability. Specifically, we showcased the inferred
 382 frequency distributions from the dLIF prior across critical EEG channels, clearly linking decreased
 383 oscillatory frequency to increased disease severity. Additionally, the inferred event-relational graphs
 384 revealed meaningful connectivity patterns that align with known neurophysiological changes in de-
 385 mentia. Lastly, boundary time prediction comparisons highlighted BayesENDS' precision in tempo-
 386 ral modeling compared to baseline methods, underscoring its superior capability to capture accurate
 387 event timings.

388 6.1 DLIF INFERRED FREQUENCY VISUALIZATION

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



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



440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

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.

6.2 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.

6.3 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. 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.

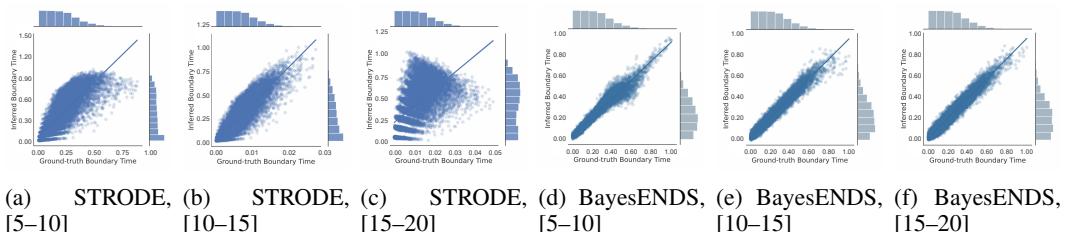


Figure 4: Predicted vs. ground-truth boundary times across frequency bands ([5–10], [10–15], [15–20] Hz): STRODE vs. BayesENDS.

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.

486 8 ETHICS STATEMENT
487488 Compliant with ICLR ethics. No human/animal subjects or IRB approval required. Public datasets
489 used under license; no re-identification attempted. For academic use; real-world applications need
490 further validation. No harmful insights, conflicts, or sponsorship. Possible dataset biases noted.
491 Methods documented for reproducibility.
492493 9 REPRODUCIBILITY STATEMENT
494495 We have taken deliberate steps to facilitate the reproducibility of our results. A complete description
496 of the model architectures, training procedures, and evaluation protocols is provided in the main
497 text. All datasets used are publicly available and cited with their licenses.
498

499 500 REFERENCES

- 501 Madhav Acharya, Ravinesh C. Deo, Xiaohui Tao, Prabal Datta Barua, Aruna Devi, Anirudh
502 Atmakuru, and Ru-San Tan. Deep learning techniques for automated Alzheimer’s and mild
503 cognitive impairment disease using EEG signals: A comprehensive review of the last decade
504 (2013–2024). *Computer Methods and Programs in Biomedicine*, 259:108506, 2025. doi:
505 10.1016/j.cmpb.2024.108506.
- 506 Hamdi Altaheri, Ghulam Muhammad, and Mansour Alsulaiman. Physics-informed attention tem-
507 poral convolutional network for eeg-based motor imagery classification. *IEEE transactions on*
508 *industrial informatics*, 19(2):2249–2258, 2022.
- 509 Claudio Babiloni et al. International federation of clinical neurophysiology (ifcn) recommendations
510 on electroencephalography in Alzheimer’s disease and other dementias. *Clinical Neurophysiology*, 2021.
- 511 André M. Bastos and Jan-Mathijs Schoffelen. A tutorial review of functional connectivity analysis
512 methods and their interpretational pitfalls. *Frontiers in Systems Neuroscience*, 9:175, 2016. doi:
513 10.3389/fnsys.2015.00175.
- 514 Guoqiang Q. Bi and Mu-Ming Poo. Synaptic modifications in cultured hippocampal neurons: depen-
515 dence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of Neuroscience*,
516 18(24):10464–10472, 1998. doi: 10.1523/JNEUROSCI.18-24-10464.1998.
- 517 Anthony N. Burkitt. A review of the integrate-and-fire neuron model: I. homogeneous synaptic
518 input. *Biological Cybernetics*, 95(1):1–19, 2006. doi: 10.1007/s00422-006-0068-6.
- 519 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary
520 differential equations. *Advances in neural information processing systems*, 31, 2018.
- 521 Justin Dauwels, François Vialatte, and Andrzej Cichocki. A comparative study of synchrony mea-
522 sures for the early diagnosis of Alzheimer’s disease based on eeg. *Clinical Neurophysiology*, 121
523 (1):81–94, 2010.
- 524 Mohsen Ehteshamzad et al. Assessing the potential of EEG in early detection of Alzheimer’s disease.
525 *Alzheimer’s & Dementia: Diagnosis, Assessment & Disease Monitoring*, 2024. doi: 10.3233/
526 ADR-230159.
- 527 Daniel E. Feldman. The spike-timing dependence of plasticity. *Neuron*, 75(4):556–571, 2012. doi:
528 10.1016/j.neuron.2012.08.001.
- 529 Mansura Habiba and Barak A Pearlmutter. Neural ordinary differential equation based recurrent
530 neural network model. In *2020 31st Irish signals and systems conference (ISSC)*, pp. 1–6. IEEE,
531 2020.
- 532 Hengguan Huang, Hongfu Liu, Hao Wang, Chang Xiao, and Ye Wang. Strode: Stochastic boundary
533 ordinary differential equation. In *International Conference on Machine Learning*, pp. 4435–4445.
534 PMLR, 2021.

- 540 Cosimo Ieracitano, Nadia Mammone, Amir Hussain, and Francesco Carlo Morabito. A novel multi-
 541 modal machine learning based approach for automatic classification of eeg recordings in demen-
 542 tia. *Neural Networks*, 123:176–190, 2020. doi: 10.1016/j.neunet.2019.12.006.
- 543
- 544 Jaeseung Jeong. Eeg dynamics in patients with alzheimer’s disease. *Clinical Neurophysiology*, 115
 545 (7):1490–1505, 2004. doi: 10.1016/j.clinph.2004.01.001.
- 546 Pramod Kachare, Digambar Puri, Sandeep B Sangle, Ibrahim Al-Shourbaji, Abdoh Jabbari,
 547 Raimund Kirner, Abdalla Alameen, Hazem Migdady, and Laith Abualigah. Lcadnet: a novel
 548 light cnn architecture for eeg-based alzheimer disease detection. *Physical and Engineering Sci-
 549 ences in Medicine*, 47(3):1037–1050, 2024.
- 550 Seul-Kee Kim, Hayom Kim, Sang Hee Kim, Jung Bin Kim, and Laehyun Kim.
 551 Electroencephalography-based classification of alzheimer’s disease spectrum during computer-
 552 based cognitive testing. *Scientific Reports*, 14, 2024. doi: 10.1038/s41598-024-55656-8.
- 553
- 554 Dominik Klepl, Min Wu, and Fei He. Graph neural network-based EEG classification: A survey.
 555 *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 32:493–503, 2024. doi:
 556 10.1109/TNSRE.2024.3355750.
- 557 Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and
 558 Brent J Lance. Eegnet: a compact convolutional neural network for eeg-based brain–computer
 559 interfaces. *Journal of neural engineering*, 15(5):056013, 2018.
- 560 Keyvan Mahjoory, Vadim V. Nikulin, Léon Bottrel, Klaus Linkenkaer-Hansen, Mario Massimo Fato,
 561 and Stefan Haufe. Consistency of EEG source localization and connectivity estimates. *NeuroIm-
 562 age*, 152:590–601, 2017. doi: 10.1016/j.neuroimage.2017.02.076.
- 563
- 564 Christoph M. Michel and Denis Brunet. EEG source imaging: A practical review of the analysis
 565 steps. *Frontiers in Neurology*, 10:325, 2019. doi: 10.3389/fneur.2019.00325.
- 566
- 567 Christoph M. Michel, Micah M. Murray, G. Lantz, Sara Gonzalez Andino, Luigi Spinelli, and
 568 R. Grave de Peralta Menendez. EEG source imaging. *Clinical Neurophysiology*, 115(10):2195–
 2222, 2004. doi: 10.1016/j.clinph.2004.06.001.
- 569
- 570 Andreas Miltiadous, Katerina D. Tzimourta, Theodora Afrantou, Panagiotis Ioannidis, Nikolaos
 571 Grigoriadis, Dimitrios G. Tsalikakis, Pantelis Angelidis, Markos G. Tsipouras, Evripidis Glavas,
 572 Nikolaos Giannakeas, and Alexandros T. Tzallas. ”a dataset of eeg recordings from: Alzheimer’s
 573 disease, frontotemporal dementia and healthy subjects”, 2024.
- 574
- 575 Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking
 576 neural networks: Bringing the power of gradient-based optimization to spiking neural networks.
IEEE Signal Processing Magazine, 36(6):61–74, 2019.
- 577
- 578 Robert L Peach, Matteo Vinao-Carl, Nir Grossman, Michael David, Emma Mallas, David Sharp,
 579 Paresh A Malhotra, Pierre Vandergheynst, and Adam Gosztolai. Implicit gaussian process repre-
 580 sentation of vector fields over arbitrary latent manifolds. *arXiv preprint arXiv:2309.16746*, 2023.
- 581
- 582 Aruane Mello Pineda, Fernando M. Ramos, Luiz Eduardo Betting, and Andriana S. L. O. Cam-
 583 panharo. Use of complex networks for the automatic detection and the diagnosis of Alzheimer’s
 584 disease. In Ignacio Rojas, Gonzalo Joya, and Andreu Catala (eds.), *Advances in Computational
 585 Intelligence: 15th International Work-Conference on Artificial Neural Networks (IWANN 2019),
 586 Proceedings, Part I*, volume 11506 of *Lecture Notes in Computer Science*, pp. 115–126, Cham,
 587 2019. Springer. doi: 10.1007/978-3-030-20521-8_10.
- 588
- 589 Seyed-Ali Sadegh-Zadeh, Elham Fakhri, Mahboobe Bahrami, Elnaz Bagheri, Razieh Khamse-
 590 hashari, Maryam Noroozian, and Amir M Hajiyavand. An approach toward artificial intelligence
 591 alzheimer’s disease diagnosis using brain signals. *diagnostics*, 13(3):477, 2023.
- 592
- 593 Md Nurul Ahad Tawhid, Siuly Siuly, Enamul Kabir, and Yan Li. Advancing alzheimer’s
 594 disease detection: a novel convolutional neural network based framework leveraging eeg
 595 data and segment length analysis. *Brain Informatics*, 12(1):13, 2025. doi: 10.
 596 1186/s40708-025-00260-3. URL <https://braininformatics.springeropen.com/articles/10.1186/s40708-025-00260-3>.

- 594 Mário L. Vicchietti, Fernando M. Ramos, Luiz E. Betting, and Andriana S. L. O. Campanharo.
 595 Computational methods of eeg signals analysis for Alzheimer’s disease classification. *Scientific*
 596 *Reports*, 13:8184, 2023. doi: 10.1038/s41598-023-32664-8.
- 597
 598 Yihe Wang, Yu Han, Haishuai Wang, and Xiang Zhang. Contrast everything: A hierarchical con-
 599 trastive framework for medical time-series. *Advances in Neural Information Processing Systems*,
 600 36:55694–55717, 2023.
- 601
 602 Yihe Wang, Nan Huang, Taida Li, Yujun Yan, and Xiang Zhang. Medformer: A multi-granularity
 603 patching transformer for medical time-series classification. *Advances in Neural Information Pro-*
 604 *cessing Systems*, 37:36314–36341, 2024a.
- 605
 606 Yihe Wang, Nadia Mammone, Darina Petrovsky, Alexandros T. Tzallas, Francesco C. Morabito, and
 607 Xiang Zhang. ADformer: A multi-granularity transformer for EEG-based Alzheimer’s disease
 608 assessment. arXiv preprint arXiv:2409.00032, 2024b.
- 609
 610 Yihe Wang, Nan Huang, Nadia Mammone, Marco Cecchi, and Xiang Zhang. Lead: Large founda-
 611 tion model for eeg-based alzheimer’s disease detection. *arXiv preprint arXiv:2502.01678*, 2025.
- 612
 613 Xiang Zhang and Lina Yao. *Deep learning for EEG-based brain–computer interfaces: representa-*
 614 *tions, 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 THEORY

Lemma 1 (Shift–stability of IVPs (Huang et al., 2021)). *Let $e > 0$ and $U \subset \mathbb{R}^n$ be open. Let $f_1, f_2 : [a - 2e, a] \rightarrow \mathbb{R}^n$ be continuously differentiable with $\|f'_1\| \leq M$ for some $M > 0$. Consider*

$$y'_1(t) = f_1(t), \quad y_1(a - e) = x_1, \quad y'_2(t) = f_2(t) = f_1(t - e), \quad y_2(a - e) = x_2.$$

Then, as $e \rightarrow 0^+$,

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

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

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

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

$$g(m) := -\frac{q(M)}{m M} \log \frac{q(M)}{M r(M) e^{-\int_0^M r(u) du}}, \quad G'(m) = g(m), \quad G(-e^{-S}) = 0. \quad (20)$$

Then

$$\text{KL}(q \| p_r) = \int_0^S q(t) \log \frac{q(t)}{r(t) e^{-\int_0^t r(u) du}} dt = \lim_{\varepsilon \downarrow 0} G(-\varepsilon), \quad (21)$$

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

$$\text{KL}(q \| p_r) \leq G(-\varepsilon) + |G(-2\varepsilon) - G(-\varepsilon)| =: \mathcal{U}_\varepsilon, \quad (22)$$

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

648 In training we evaluate the KL term using the computable bound \mathcal{U}_ε from Eq. 22 (fixed small ε and
 649 an ODE solver for Eq. (20)).
 650

651 We then analyze how entry-wise perturbations of lags affect the learned ERG when the edge map
 652 is exponential. For channels $i \neq j$ and time t , let the noise-free lag be $\Delta t_{ij}(t; T)$ and the perturbed
 653 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
 654 slope parameter $\alpha > 0$ and the (noise-free and perturbed) edges

$$655 \quad e_{ij}(t; T) = \phi_\alpha(\Delta t_{ij}(t; T)), \quad \widetilde{e}_{ij}(t; T) = \phi_\alpha(\widetilde{\Delta t}_{ij}(t; T)).$$

656 The decoder uses the Monte-Carlo, time-averaged adjacencies
 657

$$658 \quad \bar{A}_{ij} = \frac{1}{MS} \sum_{m=1}^M \sum_{s=1}^S e_{ij}(t_m; T^{(s)}), \quad \widetilde{\bar{A}}_{ij} = \frac{1}{MS} \sum_{m=1}^M \sum_{s=1}^S \widetilde{e}_{ij}(t_m; T^{(s)}).$$

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

$$663 \quad |\phi_\alpha(x) - \phi_\alpha(y)| \leq \alpha |x - y|. \quad (23)$$

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

$$666 \quad |\widetilde{e}_{ij}(t; T) - e_{ij}(t; T)| \leq \alpha |\xi_{ij}(t; T)|. \quad (24)$$

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

$$669 \quad |\widetilde{\bar{A}}_{ij} - \bar{A}_{ij}| \leq \alpha \overline{|\xi_{ij}|}, \quad \overline{|\xi_{ij}|} := \frac{1}{MS} \sum_{m=1}^M \sum_{s=1}^S |\xi_{ij}(t_m; T^{(s)})|, \quad (25)$$

672 *and the matrix (Frobenius-norm) bound*

$$674 \quad \|\widetilde{\bar{A}} - \bar{A}\|_F \leq \frac{\alpha}{MS} \sum_{m=1}^M \sum_{s=1}^S \|\Xi^{(m,s)}\|_F, \quad \Xi^{(m,s)} := [\xi_{ij}(t_m; T^{(s)})]_{i \neq j}, \quad (26)$$

677 *hence $\|\widetilde{\bar{A}} - \bar{A}\|_F \leq \alpha \|\Xi\|_F$ with $\|\Xi\|_F$ the average Frobenius norm of lag-noise matrices.*

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

$$681 \quad |\widetilde{\bar{A}}_{ij} - \bar{A}_{ij}| \leq \alpha \varepsilon_\infty, \quad \|\widetilde{\bar{A}} - \bar{A}\|_F \leq \alpha \varepsilon_\infty \sqrt{C(C-1)}. \quad (27)$$

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

$$686 \quad \mathbb{P}(|\Delta_{ij}| \geq \tau) \leq 2 \exp\left(-\frac{MS\tau^2}{2\alpha^2\sigma^2}\right). \quad (28)$$

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

$$690 \quad \mathbb{E}[|\Delta_{ij}|] \leq \alpha \sigma \sqrt{\frac{2}{\pi}}, \quad \mathbb{E}[\|\widetilde{\bar{A}} - \bar{A}\|_F] \leq \alpha \sigma \sqrt{\frac{2}{\pi}} \sqrt{C(C-1)}. \quad (29)$$

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

697 C PROOF OF THEOREM 2

699 *Proof.* By definition,

$$701 \quad \text{KL}(q\|p_r) = \int_0^S q(t) \log \frac{q(t)}{r(t)e^{-\int_0^t r(s) ds}} dt.$$

702 Using $m = -e^{-t}$, we have $t = -\log(-m) = M$ and $dt = -\frac{1}{m} dm$. Hence
 703

$$704 \text{KL}(q\|p_r) = \int_{-e^{-s}}^0 \frac{q(M)}{M} \log \frac{q(M)}{Mr(M)e^{-\int_0^M r}} \left(-\frac{1}{m} \right) dm = \int_{-e^{-s}}^0 g(m) dm.$$

707 This improper integral equals $\lim_{\varepsilon \downarrow 0} \int_{-e^{-s}}^{-\varepsilon} g(m) dm = \lim_{\varepsilon \downarrow 0} G(-\varepsilon)$, proving Eq. (21). Since
 708 $r(t) \in [a, b]$ and q is integrable, g is locally integrable near $m = 0^-$. Split the integral at $-\varepsilon$ and at
 709 -2ε :

$$710 \int_{-e^{-s}}^0 g = \int_{-e^{-s}}^{-\varepsilon} g + \int_{-\varepsilon}^0 g \leq G(-\varepsilon) + \left| \int_{-\varepsilon}^0 g \right| \leq G(-\varepsilon) + \left| \int_{-2\varepsilon}^{-\varepsilon} g \right| = G(-\varepsilon) + |G(-2\varepsilon) - G(-\varepsilon)|,$$

713 which yields Eq. (22). The last inequality uses Lemma 1 (applied to the IVPs $G'_1(m) = g(m)$
 714 and $G'_2(m) = g(m)$ shifted by ε) to control the tail near the singular endpoint and the fact that
 715 r is bounded in $[a, b]$, ensuring g remains controlled as $m \rightarrow 0^-$. As $\varepsilon \downarrow 0$, the tail vanishes by
 716 dominated convergence, hence $\mathcal{U}_\varepsilon \rightarrow \text{KL}(q\|p_r)$. \square
 717

718 D PROOF OF THEOREM 3

720 *Proof. (Lipschitzness).* The absolute value is 1-Lipschitz: $\|x| - |y\| \leq |x - y|$. The function
 721 $u \mapsto e^{-\alpha u}$ on $u \geq 0$ has derivative $|\alpha e^{-\alpha u}| \leq \alpha$, hence it is α -Lipschitz on $\mathbb{R}_{\geq 0}$. By composition
 722 of Lipschitz maps,

$$724 |\phi_\alpha(x) - \phi_\alpha(y)| = |e^{-\alpha|x|} - e^{-\alpha|y|}| \leq \alpha \|x| - |y\| \leq \alpha |x - y|,$$

725 establishing Eq. (23). Taking $y = x + \xi$ gives Eq. (24).

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

$$729 |\tilde{A}_{ij} - \bar{A}_{ij}| = \left| \frac{1}{MS} \sum_{m,s} (\tilde{e}_{ij}(t_m; T^{(s)}) - e_{ij}(t_m; T^{(s)})) \right| \leq \frac{1}{MS} \sum_{m,s} |\tilde{e}_{ij} - e_{ij}| \leq \frac{\alpha}{MS} \sum_{m,s} |\xi_{ij}(t_m; T^{(s)})|,$$

731 which is Eq. (25).

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

$$736 \|\tilde{A} - \bar{A}\|_{\text{F}} \leq \frac{1}{MS} \sum_{m,s} \|\Delta^{(m,s)}\|_{\text{F}}.$$

738 Entry-wise inequality Eq. 24 implies $\|\Delta^{(m,s)}\|_{\text{F}} \leq \alpha \|\Xi^{(m,s)}\|_{\text{F}}$, giving Eq. (26). \square
 739

740 E PROOF OF COROLLARY 4

743 *Proof.* (i) From Eq. (25), $|\xi_{ij}| \leq \varepsilon_\infty$, giving the entry-wise claim. For the matrix bound,
 744 $\|\Xi^{(m,s)}\|_{\text{F}} \leq \varepsilon_\infty \sqrt{C(C-1)}$ for every (m, s) ; apply Eq. (26).

745 (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.
 746 (24), $Y_{m,s}$ is an α -Lipschitz transform of $\xi_{ij}(t_m; T^{(s)})$, hence $Y_{m,s}$ is sub-Gaussian with proxy
 747 $\alpha^2 \sigma^2$ (standard Lipschitz preservation of the ψ_2 -norm). Then $\Delta_{ij} = (MS)^{-1} \sum_{m,s} Y_{m,s}$ is
 748 sub-Gaussian with proxy $\alpha^2 \sigma^2 / (MS)$, yielding Eq. (28) via the Chernoff bound.

750 For the Gaussian-mean bound, use $|Y_{m,s}| \leq \alpha |\xi_{ij}(t_m; T^{(s)})|$ and linearity:
 751

$$752 753 754 \mathbb{E}|\Delta_{ij}| \leq \frac{\alpha}{MS} \sum_{m,s} \mathbb{E}|\xi_{ij}| = \alpha \mathbb{E}|\xi| = \alpha \sigma \sqrt{\frac{2}{\pi}},$$

755 where the last equality is the mean absolute value of a zero-mean Gaussian. Summing the en-
 try-wise bounds in quadrature gives the Frobenius expectation in Eq. (29). \square

756 F ADDITIONAL EXPERIMENTAL DETAILS
757758 F.1 ALZHEIMER’S DISEASE EEG BASELINE METHODS
759

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

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

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

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

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

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

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

790 F.2 DATASET DESCRIPTIONS
791792 F.2.1 TOY DATASET AND DATA GENERATION
793

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

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

807 **Sequence generation.** Each synthetic sequence comprises 20 observations, constructed by sam-
808 pling inter-event times Δt_i from an exponential distribution with parameter λ . The event timestamps

810 t_i are obtained cumulatively by $t_i = \sum_{j=1}^i \Delta t_j$. Observations y_i are subsequently generated using
 811 the relationship:

$$813 \quad 814 \quad y_i = \sin(t_i) + \eta_i, \quad \eta_i \sim \mathcal{N}(0, \sigma_\eta^2),$$

816 where the default noise level is $\sigma_\eta = 0.07$. Additional sensitivity analyses varied σ_η within
 817 $\{0.05, 0.10, 0.15\}$ to assess model robustness.

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

825 F.2.2 ALZHEIMER’S DISEASE EEG DATASET

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

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

844 F.3 EXPERIMENTAL SETUP

846 We evaluate BayesENDS under a five-fold cross-subject protocol to ensure that generalization is
 847 assessed on previously unseen participants rather than unseen windows from the same individuals.
 848 For each cohort, subjects are partitioned into five disjoint folds with stratification at the subject
 849 level so that the class proportions of Alzheimer’s disease, frontotemporal dementia/mild cognitive
 850 impairment, and healthy controls are approximately preserved in every split. In each round, four
 851 folds are used for training, and one for testing; the roles of the folds are rotated until every fold
 852 serves exactly once in the held-out test set.

853 EEG is segmented into non-overlapping two-second windows (1,000 samples at 500 Hz) per subject
 854 and channel, followed by channel-wise z -normalization computed within the training portion of the
 855 active fold and then applied to validation and test windows of that fold.

856 Evaluation is conducted at subject levels. Subject-level predictions aggregate a subject’s windows
 857 by majority vote over window-wise labels. We report accuracy, macro-F1, and summarize perfor-
 858 mance as the mean and standard deviation across the five test folds.

860 G ARCHITECTURES AND TRAINING DETAILS

862 This section gives a concise description of the components used in BayesENDS and the training
 863 protocol.

864
865

G.1 INPUT, PREPROCESSING, AND WINDOWING

866 EEG is segmented into non-overlapping windows and z -scored channel-wise using statistics com-
867 puted on the training split of each fold. In our AD experiments we use 2,s windows from 500,Hz
868 recordings ($T=1000$) and $C=19$ electrodes (10–20 layout). Other datasets can adjust C and T
869 without changing the architecture.

870

871 G.2 ENCODER
872873 The encoder follows the temporal–spatial factorization popular in EEGNet, with a mild max–norm
874 constraint on spatial depthwise kernels for stability on EEG.

875

- **Block 1 (temporal → spatial).** Depthwise temporal convolution → BatchNorm → ELU; then depthwise *spatial* convolution across electrodes → BatchNorm → ELU; time average pooling; dropout (0.1).
- **Block 2 (depthwise–separable temporal).** Depthwise temporal convolution → BatchNorm → ELU; pointwise mixing → BatchNorm → ELU; time average pooling; dropout (0.1).
- **Two branches.** (a) A flattened *main* feature vector is used by the classifier; (b) a temporally downsampled, per-electrode feature map feeds the EPDE/MELP/dLIF block.

885

G.3 EPDE + MELP + DLIF COUPLING (LATENT EVENT DYNAMICS)

887 Given the encoder’s per-electrode temporal features, the EPDE produces a differentiable posterior
888 over next–event times. We parameterize a small MLP per channel to output mixture parameters for
889 the **MELP** (lognormal mixture; $K=3$ components). Sampling is reparameterized during training;
890 at test time we use mixture expectations. A compact hidden state is evolved with an explicit–Euler
891 solver to obtain a denoised per-electrode trajectory used downstream. A rate proxy read from the
892 EPDE is softly aligned with the **dLIF** prior via an L_2 rate consistency term with refractory gating
893 and a plausible alpha/theta–to–beta frequency range.

894

895 G.4 EVENT–RELATIONAL GRAPH (ERG) AND GCN
896897 From posterior samples of event times, cross–channel lags are mapped through a smooth
898 STDP–shaped nonlinearity to edge scores in $[0, 1]$, then averaged over time/samples to produce a
899 symmetric adjacency. A weak Fisher– z alignment term biases edges toward observable correlations
900 computed from the raw EEG without enforcing them. A single GCN layer converts per–channel
901 temporal descriptors into compact node embeddings that are flattened for fusion.

902

903 G.5 CLASSIFIER AND FUSION

904

905 We concatenate the encoder’s main vector with the flattened GCN features and apply a two–layer
906 MLP followed by a linear layer and Softmax over classes. No attention or recurrence is used at this
907 stage; temporal information is already summarized by EPDE/MELP and the encoder.

908

909 G.6 OPTIMIZATION AND PROTOCOL

910

- **Loss.** Cross–entropy for labels plus small auxiliary terms: EPDE/MELP reconstruction/regularizers, dLIF rate consistency, ERG Fisher– z , and the IVP–KL surrogate. We use modest default weights and found them robust across cohorts.
- **Training.** Adam ($lr 5 \times 10^{-4}$, weight decay 10^{-4}), batch size 1024, gradient–norm clipping at 1.0, 30 epochs. Learning rate is halved if validation AUC does not improve for 15 epochs; the best AUC checkpoint is kept.
- **Evaluation.** Five–fold *cross–subject* splits; subject–level predictions from window probabilities via majority vote (or simple averaging).

918 G.7 SHAPES SUMMARY
919
920921 Below we list only input/output sizes and activations for clarity; $L=250$ denotes the temporal length
922 after the first time-pooling stage.
923
924
925
926

Module	Input	Output	Activation
Input window	$\mathbb{R}^{B \times 1 \times 19 \times 1000}$	—	—
Encoder (temporal branch)	$\mathbb{R}^{B \times 1 \times 19 \times 1000}$	$\mathbb{R}^{B \times 19 \times L} (L = 250)$	ELU
Encoder (main branch, flattened)	$\mathbb{R}^{B \times 1 \times 19 \times 1000}$	$\mathbb{R}^{B \times 1178}$	ELU
EPDE/MELP/dLIF block	$\mathbb{R}^{B \times 19 \times L}$	$\mathbb{R}^{B \times 19 \times L}$	ELU (MLPs), Tanh (ODE)
ERG adjacency	lags from EPDE/MELP	$\mathbb{R}^{B \times 19 \times 19}$	exp kernel
GCN node embeddings	$\mathbb{R}^{B \times 19 \times L}$, adjacency	$\mathbb{R}^{B \times 19 \times 64}$	ReLU
Flattened graph features	$\mathbb{R}^{B \times 19 \times 64}$	$\mathbb{R}^{B \times 1216}$	—
Fusion vector	concat(main, graph)	$\mathbb{R}^{B \times 2394}$	—
Classifier logits	$\mathbb{R}^{B \times 2394}$	$\mathbb{R}^{B \times \mathcal{Y} }$	ReLU (hidden), Softmax (out)

937
938
939
940 G.8 PSEUDOCODE SUMMARIES
941
942
943
944
945
946947 **Algorithm 1** MELP Sampling (per electrode and time proxy)

948 **Require:** Mixture weights $w \in \Delta^{K-1}$, means $\mu \in \mathbb{R}^K$, standard deviations $\sigma \in \mathbb{R}_+^K$
 949 **Ensure:** Inter-event interval $\tau \in \mathbb{R}_+$
 950 1: Sample component $k \sim \text{Categorical}(w)$
 951 2: Sample noise $\epsilon \sim \mathcal{N}(0, 1)$
 952 3: $\tau \leftarrow \exp(\mu_k + \sigma_k \cdot \epsilon)$
 953 4: **return** τ

955
956
957
958
959
960
961 **Algorithm 2** Neural ODE evolution over $[t_s, t_e]$ with S Euler sub-steps
962

963 **Require:** Features ξ , projection map $\text{proj}(\cdot)$, decoder $\text{decode}(\cdot)$, vector field $f(\cdot)$, residual weight
 964 $\alpha > 0$, sub-steps $S \in \mathbb{N}$, interval $[t_s, t_e]$
 965 **Ensure:** Decoded trajectory $\hat{z} \in \mathbb{R}^{\lfloor T/4 \rfloor}$
 966 1: $y_0 \leftarrow \text{proj}(\xi)$
 967 2: $\Delta t \leftarrow (t_e - t_s)/S$
 968 3: **for** $m = 0, 1, \dots, S-1$ **do**
 969 4: $y_{m+1} \leftarrow y_m + \Delta t \cdot (f(y_m) + \alpha y_m)$
 970 5: **end for**
 6: $\hat{z} \leftarrow \text{decode}(y_S)$
 7: **return** \hat{z}

972 **Algorithm 3** Graph weights from z_t

973

974 **Require:** Per-electrode trajectories $z_t \in \mathbb{R}^{N \times L}$ with $L = \lfloor T/4 \rfloor$; scale $\gamma > 0$; kernel mode

975 mode $\in \{\text{exp, gauss, inv1}\}$

976 **Ensure:** Symmetric adjacency $W \in \mathbb{R}^{N \times N}$ with zero diagonal

977 1: $\forall c \in \{1, \dots, N\} : s_c \leftarrow \frac{1}{L} \sum_{t=1}^L z_t(c, t)$

978 2: **for** $i = 1, \dots, N$ **do**

979 3: **for** $j = 1, \dots, N$ **do**

980 4: **if** $i = j$ **then**

981 5: $W_{ij} \leftarrow 0$

982 6: **else**

983 7: $\Delta \leftarrow |s_i - s_j|$

984 8: $W_{ij} \leftarrow \exp(-\gamma \Delta)$

985 9: **end if**

986 10: **end for**

987 11: **end for**

988 12: $W \leftarrow \frac{1}{2}(W + W^\top)$ ▷ Enforce symmetry

13: **return** W

Algorithm 4 BayesENDS: Training and Inference

Require: Dataset $\mathcal{D} = \{(X, Y)\}$, electrodes C , window length T , MELP comps K , ODE substeps S

Require: Loss weights $\lambda_{\text{aux}}, \lambda_{\text{spk}}, \lambda_{\text{graph}}, \lambda_{\text{LIF}}$; LR η

Ensure: Trained params Θ ; predictor BayesENDS(\cdot)

- 1: Initialize encoder ψ , EPDE+MELP ϕ , dLIF head ξ , ERG+GCN η_g , classifier θ ; Adam(η)
- 2: **for** epoch = 1.. E **do**
- 3: **for** mini-batch $(X, Y) \sim \mathcal{D}$ **do**
- 4: **Preprocess:** channel-wise z -score per window
- 5: **Encoder** (ψ): $\text{main_vec} \in \mathbb{R}^{B \times d_{\text{main}}}$, $\text{temp_feat} \in \mathbb{R}^{B \times C \times 1 \times L}$, $L = \lfloor T/4 \rfloor$
- 6: **EPDE+MELP+ODE:**
For $c = 1..C$: EPDE \Rightarrow mixture $(\mathbf{w}_c, \boldsymbol{\mu}_c, \boldsymbol{\sigma}_c)$; sample $\tau = \exp(\mu_{c,k} + \sigma_{c,k}\varepsilon)$, accumulate events $\{T_c\}$; ODE evolve with S Euler steps \Rightarrow trajectory z_c :
 - 7: Stack $Z \in \mathbb{R}^{B \times C \times L}$; keep event lags for ERG
 - 8: **dLIF prior:** compute $r_c(t)$ from Z with refractory gate; rate proxy $\hat{r}_c(t)$
 - 9: **Regularizers:** $\mathcal{R}_{\text{LIF}} = \sum_c \int (\hat{r}_c - r_c)^2$; $\text{KL}_T = \sum_c \mathcal{U}_\varepsilon(q_c \| p_{\text{dLIF}})$ (IVP-KL bound)
 - 10: **ERG** from event lags: $e_{ij}(t) = \exp(-\alpha |\Delta \tilde{t}_{ij}(t)|)$, average $\Rightarrow \bar{A}$
 - 11: **Fisher-z alignment:** $\mathcal{R}_{\text{ERG}} = \sum_{i < j} [(z_{ij}^{\text{obs}} - z_{ij}^{\text{pred}})^2 / (2\sigma_{ij}^2) + \frac{1}{2} \log \sigma_{ij}^2]$
 - 12: **GCN** on (\bar{A}, Z) : $G \in \mathbb{R}^{B \times (C \cdot d_g)}$; **Fuse** $H = [\text{main_vec}; G]$
 - 13: **Classifier:** logits $\ell = \text{MLP}(H)$, $p = \text{softmax}(\ell)$
 - 14: **Loss:** $\mathcal{L} = \text{CE}(p, Y) + \lambda_{\text{aux}} \mathcal{L}_{\text{STRODE}} + \lambda_{\text{spk}} \mathcal{L}_{\text{spk}} + \lambda_{\text{graph}} \mathcal{R}_{\text{ERG}} + \lambda_{\text{LIF}} \mathcal{R}_{\text{LIF}} + \text{KL}_T$
 - 15: Backprop; clip $\|\nabla\| \leq 1$; Adam step on Θ
 - 16: **end for**
 - 17: **end for**
- 18: **Inference on window X :**
- 19: Preprocess \rightarrow Encoder; EPDE gives MELP expectations \rightarrow events $\{T_c\}$; ODE $\rightarrow Z$; build \bar{A} ; GCN $\rightarrow G$; fuse $\rightarrow H$; output p and $\hat{y} = \arg \max p$
- 20: **Subject aggregation:** majority vote $\Rightarrow \hat{y}_{\text{subj}}$