ispiacademy : Pages:1–18

Original Research

Causal-Consistent Budgeted State-Space Identification for Cross-Modal Seizure Forecasting and Closed-Loop Intervention

Suresh Bhatta[1]

[1]Lumbini Buddhist University, Faculty of Science and Technology, Tansen Road, Butwal 32907, Nepal.

Abstract

Epileptic seizures reflect transient transitions in neural population dynamics that can unfold over seconds to minutes, yet practical monitoring and therapy selection must operate under partial observability, sensor heterogeneity, and strict computational limits. Wearable scalp EEG offers long-horizon coverage but coarse spatial resolution and frequent artifacts, while intracranial recordings provide localized access to propagating activity but only for short clinical windows and with patient-specific sampling geometries. These constraints motivate algorithms that are not only accurate, but also resource-aware, uncertainty-calibrated, and consistent across sensing modalities. This paper proposes a unified technical framework for seizure forecasting and intervention planning based on causal-consistent, budgeted identification of latent state-space models with directed interaction structure. The core contribution is a learning and inference principle that enforces invariances across modalities and recording sessions via counterfactual consistency regularization, while simultaneously optimizing a constrained objective that accounts for energy, latency, and memory budgets. The resulting model yields a belief state over seizure regimes and a sparse, time-varying directed graph that supports both low-cost wearable inference and high-resolution intracranial interpretation. We further introduce a safety-aware action selection mechanism for closed-loop neurostimulation that operates in belief space and incorporates conservative uncertainty penalties to reduce the risk of inappropriate intervention. The framework is designed to support principled personalization, robust operation under artifacts and missing channels, and explicit trade-offs between predictive performance and computational footprint.

1. Introduction

Seizure prediction and seizure-responsive therapy are often discussed as separate problems, with distinct algorithmic toolchains and success criteria [1]. Prediction emphasizes early warning, low false alarm rates, and stable performance under daily-life variability, whereas therapy planning emphasizes localization, network interpretation, and choosing targets or parameters that meaningfully change outcomes. In practice, however, these two tasks are coupled by the same underlying limitation: neural dynamics are only partially observed, and the observation mechanism changes with electrode montage, patient physiology, and context. A predictor that is optimized purely for classification can rely on superficial correlates that do not generalize across sessions or devices, while a network analysis that is optimized purely for interpretability can be too complex or fragile to deploy continuously. A more coherent approach treats both prediction and intervention as downstream uses of a single probabilistic belief state that is updated online under explicit resource constraints and that admits refinement when higher-fidelity data are available.

From a modeling standpoint, seizures can be viewed as dynamical regime transitions in a high-dimensional, nonlinear system with latent variables that are not directly measurable [2]. The observed EEG or SEEG signals conflate multiple neural sources, non-neural artifacts, and measurement effects. Moreover, seizure evolution is not static; connectivity patterns and spectral properties can vary with

sleep, medication, hormonal cycles, or electrode impedance drift. These considerations argue against purely static feature-based models and motivate state-space formulations that explicitly represent latent trajectories and their uncertainties. Yet classical state-space identification assumes stable observation operators and sufficient compute for filtering and smoothing, which conflicts with wearable constraints and with intracranial sampling heterogeneity.

This paper argues that two requirements must be treated as first-class design principles if one aims to bridge continuous monitoring and intervention planning [3]. The first requirement is budgeted inference: the algorithm must explicitly represent and optimize the computational footprint, rather than treating computational simplification as a post hoc pruning step. The second requirement is causal consistency across modalities and sessions: the learned representation must preserve predictive relationships that remain valid under changes in sensors, montage, and recording context, and must avoid spurious dependencies that arise from modality-specific artifacts or mixing effects. These requirements lead to an identification problem that differs from conventional supervised learning. The goal is not only to minimize prediction error, but also to enforce invariances that approximate counterfactual stability under changes in measurement and to constrain the solution to a feasible compute region.

Table 1: Core variables in the state-space seizure model.

Symbol	Type	Description
$y_t^{(m)}$	Raw signal	Multichannel EEG/SEEG for modality m
$x_t^{(m)}$	Features	Windowed features $\phi^{(m)}(y_{t-w+1:t}^{(m)})$
$s_t \in \mathbb{R}^N$	Latent state	Neural population state over N nodes
$r_t \in \{1, \ldots, R\}$	Regime index	Interictal, preictal, ictal, transition
$A_t \in \mathbb{R}^{N \times N}$	Directed graph	Time-varying adjacency of interactions
u_t	Action	Exogenous input, e.g., stimulation command
b_t	Belief	Posterior over (s_t, r_t, A_t) at time t

Table 2: Sensing configurations and their main characteristics.

Configuration	Coverage	Resolution	Constraints
Wearable EEG	Weeks–months	Coarse spatial, dense time	Strict energy, memory, latency
Intracranial SEEG	Days	High spatial, high SNR	Limited duration, patient-specific grid
Simultaneous EEG+SEEG	Hours	Cross-modal	Calibration and alignment windows
Ambulatory low-power	Continuous	Minimal features	Always-on, hard compute budget
Clinical high-res	Inpatient	Rich interaction features	Relaxed budget, offline analysis

Table 3: Hybrid latent representation used for seizure dynamics.

Component	Variable	Role	Example instantiation
Regime process	r_t	Seizure phase index	Interictal / preictal / ictal
Continuous state	s_t	Low-dim dynamics	Band-limited amplitudes, excitability
Directed graph	A_t	Interaction structure	Propagation and hemisphere flow
Artifact term	$a_t^{(m)}$	Non-neural effects	Motion, impedance jumps, stim artifact
Patient params	$\theta^{(p)}$	Personalization	Transition biases, noise scales

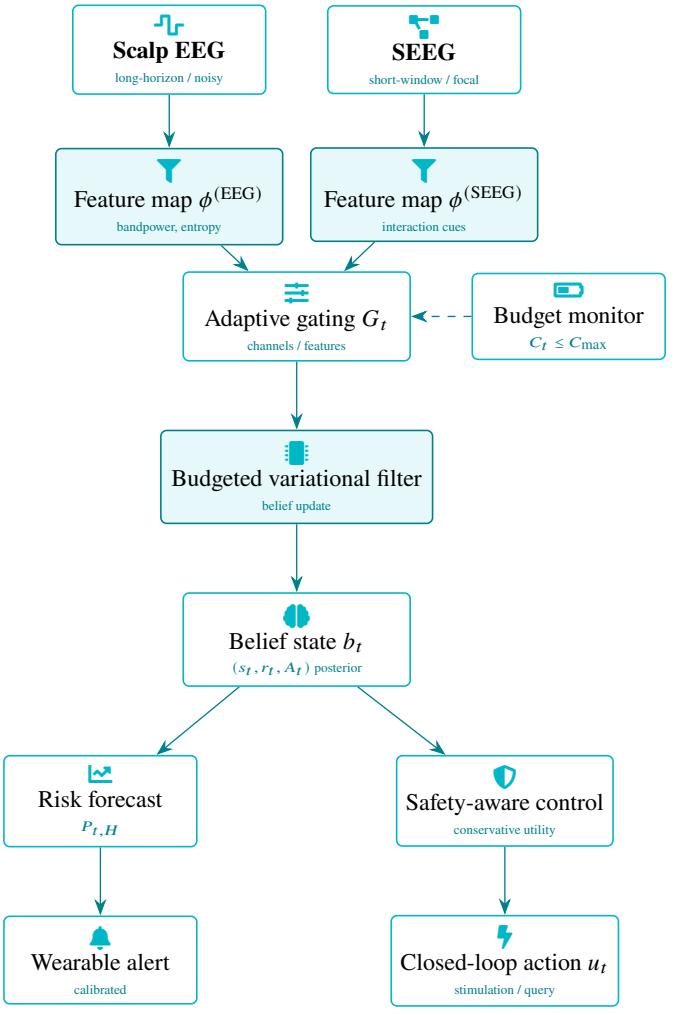

Figure 1: End-to-end cross-modal pipeline with budgeted conditional computation. Scalp EEG and SEEG features feed a gated, resource-aware variational filter that maintains a belief state over latent dynamics, regimes, and directed interactions for forecasting and conservative closed-loop intervention.

Empirically, the importance of computation-aware personalization is illustrated by reported monitoring approaches that achieve strong performance using only a small number of channels and features, emphasizing that per-patient selection can support low complexity without collapsing accuracy [4]. Separately, intracranial analyses have emphasized directed interaction structure and hemisphere-dependent network effects for informing stimulation targets and anticipating outcomes [5]. Such observations motivate a unified formulation in which low-cost wearable inference estimates the belief state and its uncertainty, while higher-resolution intracranial data refine the directed interaction structure used for planning and control, without forcing either modality to be the canonical ground truth.

The technical contributions of this paper are threefold. First, we introduce a causal-consistent identification objective that combines a latent state-space model with a sparse, time-varying directed interaction graph, and we add counterfactual consistency regularization terms that encourage invariance of predictive relationships under changes in observation operators and feature extraction. Second, we develop a budgeted variational filtering method that integrates feature gating, channel gating, and structured sparsification into the inference objective through explicit compute constraints, yielding a family of deployable

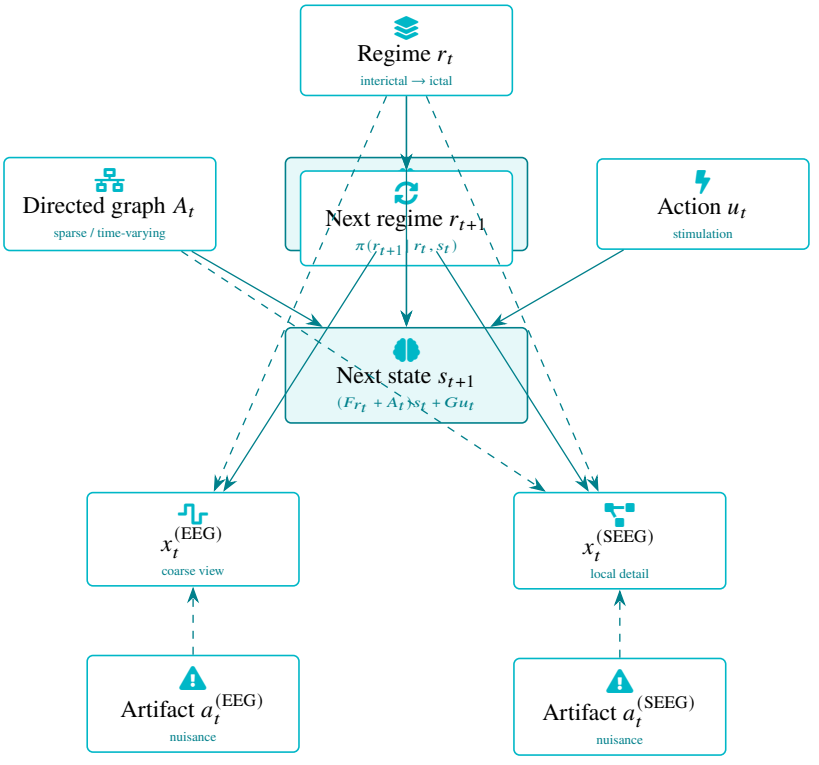

Figure 2: Hybrid latent dynamics with regime switching, directed interactions, and modality-specific observations. A single belief representation links state transitions, optional control inputs, and feature likelihoods across scalp and intracranial sensing.

Table 4: Breakdown of computational cost in budgeted inference.

Source	Cost term	Control handle	Typical tuning
Feature eval	$\sum g_{t,j}^{(m)} c_j^{(m)}$	Feature and channel gates $g_{t,j}^{(m)}$	Turn on near transitions
Latent update	$\alpha_1 d_t^2$	Effective dimension d_t	Low-rank + diagonal
Graph ops	$\alpha_2 \|A_t\|_0^{\mathrm{eff}}$	Sparsity of A_t	Group-sparse penalties
Forecasting	Horizon steps H	Rollout depth	Short horizon on-device
Control	Action scoring calls	Action set size	Few discrete templates

approximations tuned to device capabilities. Third, we propose a belief-space, safety-aware intervention selection rule for closed-loop stimulation that accounts for posterior uncertainty and imposes conservative penalties to limit risk under model mismatch and artifact contamination [6]. Throughout, personalization is treated as Bayesian adaptation of patient-specific parameters under stability constraints, rather than as a purely discriminative re-training process.

2. Problem Formulation

We consider multichannel electrophysiological observations collected from one or more modalities, such as scalp EEG and intracranial SEEG. Let $y_t^{(m)} \in \mathbb{R}^{M_m}$ denote the observed signal at discrete time t for modality m, where M_m is the number of channels. The sampling rate and preprocessing can differ across modalities, so t should be interpreted as indexing a common analysis grid, such as

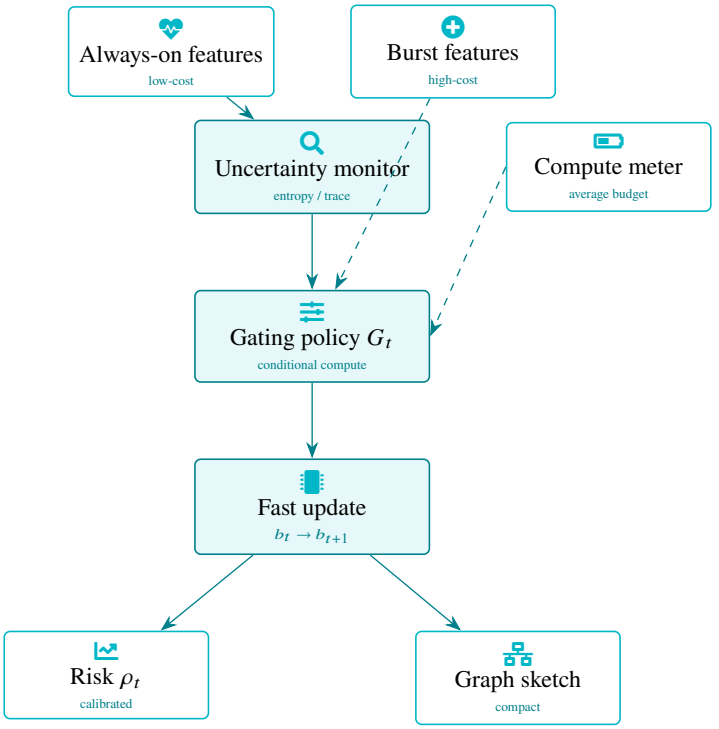

Figure 3: Adaptive gating for budgeted online inference. Baseline features run continuously, while additional computations are activated only when belief uncertainty rises, producing risk and graph summaries under strict average-cost constraints.

Table 5: Regularization and constraints in the identification objective.

Term	Target	Data region	Weight
$\mathcal{R}_{\text{post}}$	Cross-modal posterior match	Overlapping EEG/SEEG windows	γ_1
$\mathcal{R}_{\text{cf}}$	Risk invariance under perturbations	Pseudo-counterfactual inputs	γ_2
$\mathcal{R}_{\text{sparse}}$	Sparse and stable graphs	A_t, gates	γ_3
$\beta C(\lambda)$	Compute-awareness	All time steps	β
KL trust-region	Conservative adaptation	$\theta^{(p)}$ updates	Radius ϵ

feature windows of fixed duration. For each modality m, we define a feature extraction map $\phi^{(m)}$ that produces $x_t^{(m)} = \phi^{(m)}(y_{t-w+1:t}^{(m)}) \in \mathbb{R}^{D_m}$, where w is the window length in samples and D_m is the feature dimension. Feature extraction may include spectral power, entropy-like measures, wavelet coefficients, or low-rank projections. The framework does not require that $\phi^{(m)}$ be differentiable, but it must be computable under the platform's resource budget, and the choice of which features to compute is itself part of the inference problem.

The latent neural dynamics are represented by a state $s_t \in \mathbb{R}^N$ defined over spatial nodes that can be interpreted at a chosen resolution, such as cortical parcels or electrode-contact clusters. We also introduce a discrete regime variable $r_t \in \{1, \ldots, R\}$ that captures seizure-related phases, including interictal and ictal regimes, as well as intermediate transition regimes when needed. The pair (s_t, r_t) forms a hybrid latent state. To represent directed interactions among nodes, we define a time-varying adjacency matrix $A_t \in \mathbb{R}^{N \times N}$ whose entry $A_{t,ij}$ encodes the directed influence of node j on node i at

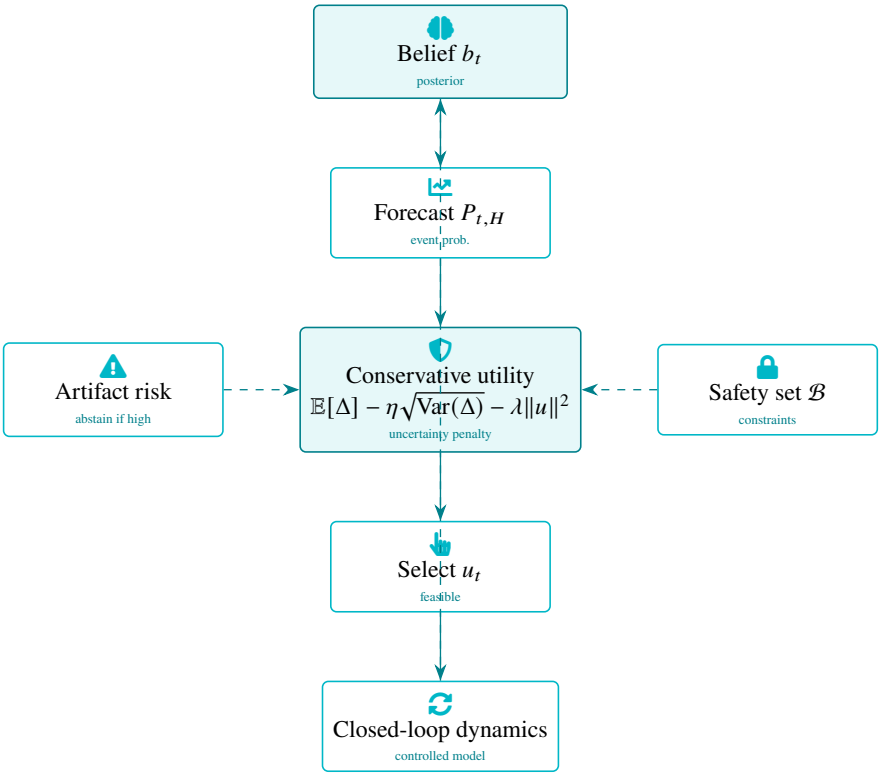

Figure 4: Safety-aware belief-space decision loop. Forecasts derived from the belief inform a conservative utility for action choice, while explicit constraints and artifact-aware abstention reduce inappropriate intervention in uncertain conditions.

Table 6: Representative operating modes of the learned model.

Mode	Inputs	Budget level	Primary function
Wearable monitoring	Scalp EEG	Lowest	Long-horizon risk tracking
Wearable + cloud	EEG + server	Low–medium	Periodic recalibration
Intracranial refinement	SEEG	High	Target and pathway mapping
Retrospective analysis	Multimodal archive	Flexible	Protocol and model design
Simulation mode	Synthetic data	Offline	Ablation and stress testing

time t. The adjacency is not assumed to be static; it can change with regime and evolve gradually within a regime [7].

A generative state transition model is given by

$$s_{t+1} = F_{r_t} s_t + A_t s_t + G u_t + \xi_t, \quad \xi_t \sim \mathcal{N}(0, \Sigma_{r_t}), \tag{2.1}$$

$$p(r_{t+1} \mid r_t, s_t) = \mathrm{Cat}\big(\pi_{r_t}(\psi, s_t)\big), \tag{2.2}$$

where F_{r_t} captures regime-dependent local dynamics, u_t denotes exogenous inputs including potential stimulation actions, and G maps actions to state effects. The regime transition probabilities $\pi_{r_t}(\psi, s_t)$ may depend on parameters ψ and on the current continuous state. For computational tractability, one may restrict π to depend on a low-dimensional summary of s_t, such as energy in certain bands or a scalar excitability coordinate.

Table 7: Forecasting outputs and evaluation metrics.

Aspect	Metric	Purpose
Discrimination	Sensitivity vs FAR	Event-level separability
Calibration	Brier score, NLL	Match of risk to event frequency
Timing	Time-to-detection	Lead time before onset
Clustering	Event-wise precision/recall	Handling of clustered seizures
Resource use	Avg / tail latency	Real-time feasibility under budget
Robustness	Shifted-montage tests	Sensitivity to modality changes

Table 8: Components of the safety-aware belief-space control objective.

Term	Argument	Role	Effect on policy
$\ell_{\text{seiz}}(r_t)$	Regime r_t	Penalize ictal time	Drives risk reduction
$\lambda_{\text{stim}}\|u_t\|^2$	Action u_t	Penalize energy	Limits intensity and duty
$\lambda_{\text{risk}}\mathcal{S}$	(s_t, A_t, u_t)	Encode safety rules	Avoids unsafe regions
$\eta\sqrt{\text{Var}(\Delta)}$	Belief b_t	Uncertainty penalty	Encourages conservatism
$\lambda_{\text{bar}}\varphi$	(b_t, u_t)	Belief barrier	Keeps beliefs in safe set

Table 9: Personalization and deployment-related parameter groups.

Group	Examples	Timescale	Constraint
Population dynamics	Shared F_r, base transitions	Pre-training	Strong priors, slow change
Patient regime biases	Transition tendencies	Days–weeks	KL-bounded updates
Observation model	Gains, noise, artifact rates	Hours–days	Regularized toward template
Device budget	$C_{\max}$, deadlines	Fixed per device	Set by hardware
Safety thresholds	Risk and entropy limits	Clinician-defined	Conservative margins

The observation model links modality-specific features to the latent state. Because each modality observes different projections of neural activity, we define modality-dependent observation operators $H^{(m)}$ and include nuisance terms for artifacts. A general form is

$$x_t^{(m)} = h^{(m)}(s_t, A_t, r_t) + a_t^{(m)} + \epsilon_t^{(m)}, \quad \epsilon_t^{(m)} \sim \mathcal{N}(0, R_t^{(m)}), \tag{2.3}$$

where $h^{(m)}$ maps latent dynamics to expected features, $a_t^{(m)}$ captures artifacts and non-neural contamination, and $R_t^{(m)}$ is a noise covariance that may be time-varying to accommodate nonstationary measurement quality. In wearable contexts, $h^{(m)}$ is often simplified to depend primarily on local summaries of s_t, whereas in intracranial contexts it can incorporate interaction-dependent features that are sensitive to directed connectivity.

The inference goal is to compute a posterior belief over (s_t, A_t, r_t) given observed features $x_{1:t}^{(m)}$ across available modalities. The forecasting goal is to estimate the probability of entering an ictal regime within a horizon H, namely $p(r_{t+\tau} \in \mathcal{I} \text{ for some } \tau \le H \mid \mathcal{D}_{1:t})$ where $\mathcal{I}$ denotes ictal regimes and $\mathcal{D}_{1:t}$ denotes all available data up to time t. The intervention goal is to choose actions u_t that reduce expected seizure burden while managing stimulation cost and safety risk, conditioned on the current belief state [8].

Two constraints shape the problem. The first is a compute budget constraint. Let C_t denote the computational cost incurred at time t, including feature computation, filtering updates, and any graph-related operations. We require $\mathbb{E}[C_t] \le C_{\max}$ for a device-specific budget $C_{\max}$. The second is a cross-modal

consistency constraint. The model must produce representations that are stable under changes in modality, montage, and session, in the sense that predictive relationships between latent variables and seizure regimes should not change drastically when the observation operator changes. This is not a claim of perfect invariance, but a practical regularization target that reduces reliance on modality-specific spurious correlates [9].

We also incorporate personalization. Let θ denote population-level parameters and $\theta^{(p)}$ denote patient-specific parameters for patient p. Patient-specific parameters include regime transition tendencies, typical adjacency sparsity patterns, and observation noise scales. These parameters are adapted online or offline with stability constraints to avoid drift. The combination of hybrid latent dynamics, time-varying directed graphs, compute constraints, and invariance regularization yields an identification and inference problem that is substantially more structured than conventional seizure classification, and it is designed to provide a coherent belief state usable for both forecasting and control.

3. Causal-Consistent Budgeted Identification Method

This section introduces the proposed identification objective and the associated inference mechanisms [10]. The guiding idea is to learn a latent model that is both resource-feasible and stable across counterfactual changes in measurement, approximated through regularization terms that enforce cross-modal and cross-session consistency of predictive relationships. The method is designed to be deployable in multiple operating modes, ranging from low-cost wearable filtering to high-resolution intracranial refinement, without changing the semantics of the latent variables.

We begin with a variational formulation. Let $\mathcal{Z}_{1:T} = \{s_{1:T}, r_{1:T}, A_{1:T}, a_{1:T}\}$ denote latent variables over a training interval. Given observed features $\mathcal{X} = \{x_t^{(m)}\}$, we seek an approximate posterior $q_\lambda(\mathcal{Z}_{1:T} \mid \mathcal{X})$ within a structured family parameterized by λ. The baseline objective maximizes an evidence lower bound:

$$\mathcal{L}(\lambda, \Theta) = \mathbb{E}_{q_\lambda}\big[\log p_\Theta(\mathcal{X}, \mathcal{Z}_{1:T}) - \log q_\lambda(\mathcal{Z}_{1:T} \mid \mathcal{X})\big], \tag{3.1}$$

where Θ collects model parameters. For sequential deployment, q_λ is chosen to factorize in a filtering form, so that online updates are feasible and smoothing is optional [11].

To incorporate compute constraints, we define a differentiable surrogate cost $C(\lambda)$ representing expected per-step compute. A practical cost model decomposes into feature costs, channel costs, and latent update costs. Let $g_{t,j}^{(m)} \in [0,1]$ denote a gate indicating whether feature j of modality m is computed at time t, and let $c_j^{(m)}$ denote its cost. Let d_t denote the effective latent update dimension after structured sparsification and low-rank approximations. A generic cost surrogate is

$$C(\lambda) = \frac{1}{T}\sum_{t=1}^{T}\Big(\sum_m \sum_j g_{t,j}^{(m)} c_j^{(m)} + \alpha_1 d_t^2 + \alpha_2 \|A_t\|_0^{\text{eff}}\Big), \tag{3.2}$$

where $\|A_t\|_0^{\text{eff}}$ is an effective sparsity measure induced by the chosen parameterization, and α_1, α_2 map algorithmic operations to cost units. The exact form can be adapted to hardware, but the key is that C is optimized jointly with statistical fit.

We then form a constrained or penalized objective:

$$\max_{\lambda,\Theta} \ \mathcal{L}(\lambda, \Theta) \quad \text{subject to} \quad C(\lambda) \le C_{\max}, \tag{3.3}$$

or equivalently maximize $\mathcal{L}(\lambda, \Theta) - \beta C(\lambda)$ for a Lagrange multiplier β. This induces feature gating and graph simplification as part of model learning, rather than as external pruning [12].

Compute constraints alone do not ensure stability across modalities. To address cross-modal consistency, we introduce counterfactual consistency regularization. The conceptual requirement is that if two

modalities observe the same underlying latent dynamics through different observation operators, then the inferred latent state distributions and their predictive relationships to seizure regimes should be consistent up to known transformations. Since true counterfactual modalities are not observed simultaneously in general, we approximate this principle through two types of regularization.

The first regularization aligns latent posteriors across modalities when data overlap in time, such as during simultaneous EEG and SEEG recording windows or during controlled calibration sessions [13]. For overlapping time indices $\mathcal{T}_{\text{ov}}$, we add

$$\mathcal{R}_{\text{post}} = \frac{1}{|\mathcal{T}_{\text{ov}}|} \sum_{t \in \mathcal{T}_{\text{ov}}} \text{KL}\Big(q_\lambda(s_t, r_t \mid \mathcal{X}^{(\text{EEG})}) \,\Big\|\, q_\lambda(s_t, r_t \mid \mathcal{X}^{(\text{SEEG})}) \Big), \tag{3.4}$$

where $\mathcal{X}^{(\text{EEG})}$ and $\mathcal{X}^{(\text{SEEG})}$ denote modality-specific feature sets. The KL direction can be symmetrized if desired, but the essential effect is to discourage modality-specific posterior shifts that are not supported by the generative model.

The second regularization enforces invariance of predictive conditionals under changes in observation maps, even when overlap is limited. We define a seizure-risk functional $q_\lambda(\rho_t \mid \mathcal{X})$, where ρ_t is a scalar representing predicted near-term seizure risk derived from the belief over regimes. We then create pseudo-counterfactual feature sets by applying stochastic perturbations that emulate montage changes, channel dropout, and feature dropout, and we penalize changes in ρ_t under these perturbations. Let $\tilde{\mathcal{X}}$ be a perturbed version of $\mathcal{X}$. We add

$$\mathcal{R}_{\text{cf}} = \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\tilde{\mathcal{X}} \sim \mathcal{P}(\mathcal{X})} \Big[\big(\mathbb{E}_{q_\lambda}[\rho_t \mid \mathcal{X}] - \mathbb{E}_{q_\lambda}[\rho_t \mid \tilde{\mathcal{X}}] \big)^2 \Big], \tag{3.5}$$

where $\mathcal{P}$ defines perturbations consistent with plausible sensor changes. This encourages the model to base risk estimates on stable latent structure rather than fragile measurement idiosyncrasies [14].

The full training objective becomes

$$\max_{\lambda, \Theta} \; \mathcal{L}(\lambda, \Theta) - \beta C(\lambda) - \gamma_1 \mathcal{R}_{\text{post}} - \gamma_2 \mathcal{R}_{\text{cf}} - \gamma_3 \mathcal{R}_{\text{sparse}}, \tag{3.6}$$

where $\mathcal{R}_{\text{sparse}}$ enforces structured sparsity on A_t and on gating variables, and $\gamma_1, \gamma_2, \gamma_3$ control regularization strength. The sparsity regularizer is designed to prefer stable, interpretable graphs and to keep inference costs bounded. A useful choice is a group-sparse penalty that encourages consistent edge sets across time within a regime while allowing regime-dependent rewiring.

To parameterize A_t efficiently, we use a low-dimensional latent representation. Let $A_t = U\text{diag}(\theta_t)V^\top + S_t$, where $U, V \in \mathbb{R}^{N \times r}$ are learned factors, $\theta_t \in \mathbb{R}^r$ are time-varying weights, and S_t is a sparse correction. This representation supports efficient multiplication and encourages shared interaction motifs across time. The sparse correction captures patient-specific idiosyncrasies and localized propagations that are not captured by the global factors [15]. The correction can be constrained to be stable within regimes via penalties on $\|S_{t+1} - S_t\|_1$ or via regime-conditioned priors.

In intracranial refinement mode, one can augment the model with node-level directional flow metrics derived from the adjacency and latent activity. Specifically, define net outgoing and incoming influence at node i as functionals of A_t and s_t, and use them to identify candidate nodes that dominate propagation in particular regimes. Related SEEG analysis approaches have used graph-derived outflux and influx metrics combined with clustering to infer minimally invasive treatment targets from selected contacts, with retrospective consistency between computational outputs and clinical decisions in focal seizure cases [16]. In our framework, such metrics arise as decision-relevant summaries of the inferred latent graph and can be computed at high resolution when SEEG data are available, while wearable operation can rely on coarser summaries without attempting fine localization.

Personalization is implemented by splitting parameters into shared and patient-specific components. Shared parameters capture general properties of seizure regimes and dynamics motifs, while patient-specific parameters capture electrode geometry effects, baseline spectral characteristics, and idiosyncratic propagation patterns [17]. Patient adaptation is constrained by trust-region updates that limit KL divergence between successive parameter posteriors, reducing drift. A convenient formulation is to update patient parameters $\Theta^{(p)}$ by maximizing a local objective subject to $\mathrm{KL}(q(\Theta^{(p)}) \| q_{\mathrm{prev}}(\Theta^{(p)})) \leq \epsilon$, which can be implemented via dual variables and yields conservative adaptation.

The budgeted identification method thus produces a family of models indexed by compute budgets and regularization strengths. In a low-budget configuration, gating variables select a small set of features and channels, r is small, and sparse corrections are limited, yielding fast filtering. In a high-budget configuration, gating opens up richer interaction features and more graph capacity, enabling finer characterization and improved target planning. Crucially, the invariance regularizers encourage both configurations to represent the same latent semantics, facilitating transfer between wearable monitoring and clinical planning [18].

4. Online Forecasting and Safety-Aware Closed-Loop Intervention

This section describes how the learned model is deployed online for forecasting and for action selection under safety constraints. The key objects are the belief state, the risk forecast, and the action-value estimate, all of which must be computed under tight latency and must degrade gracefully under artifacts and missing data.

Let b_t denote the belief state at time t, defined as the approximate posterior over (s_t, r_t, A_t) given data up to time t. The belief is updated sequentially using a filtering recursion, which can be implemented in several ways depending on compute budget. In a lightweight configuration, the filter maintains a Gaussian approximation for s_t and categorical probabilities for r_t, with A_t represented through a small set of parameters θ_t with a diagonal or low-rank covariance [19]. The update consists of a prediction step using the state transition model and an update step using the observation likelihood for the computed features.

Because the observation can be gated, the update step must explicitly account for which features are available. Let G_t denote the set of active features at time t. The likelihood is then $p(x_t^{G_t} \mid s_t, A_t, r_t)$, and the filter update uses only those components. The gating policy can be myopic or belief-driven. A belief-driven gating policy computes a scalar uncertainty measure such as posterior entropy over regimes or trace of the state covariance, and activates additional features when uncertainty exceeds a threshold, subject to budget. This allows the system to run mostly in a low-cost mode and to temporarily increase computational effort near suspected transitions [20].

Forecasting is based on propagating the belief forward over a horizon H under the transition model, producing a forecast distribution over regimes. Define the near-term seizure event $E_{t,H}$ as entering an ictal regime within H steps. The forecast probability is

$$P_{t,H} = \mathbb{P}(E_{t,H} \mid b_t). \tag{4.1}$$

Computationally, $P_{t,H}$ can be approximated by simulating a small number of belief rollouts, by linearizing regime transitions, or by using an analytically tractable approximation when r_t is modeled as a Markov chain with transition rates modulated by a low-dimensional summary of s_t. The choice depends on budget. Importantly, the forecast should be calibrated, because intervention decisions are sensitive to overconfidence. The counterfactual consistency regularizers introduced earlier serve as training-time mechanisms to reduce systematic miscalibration induced by modality-specific artifacts [21].

Closed-loop intervention introduces an action u_t that can influence the latent dynamics. The action space depends on device capabilities. In a responsive stimulation device, actions may include selecting among candidate stimulation contacts and choosing amplitude or pulse parameters. In a wearable

advisory system, actions may be limited to alerts or to requests for additional measurement. The decision objective must trade seizure suppression against stimulation cost and safety risk [22]. We define an instantaneous loss

$$\ell(s_t, r_t, u_t) = \ell_{\text{seiz}}(r_t) + \lambda_{\text{stim}}\|u_t\|^2 + \lambda_{\text{risk}}\,\mathcal{S}(s_t, A_t, u_t), \tag{4.2}$$

where $\ell_{\text{seiz}}(r_t)$ penalizes ictal regimes, $\|u_t\|^2$ penalizes stimulation energy, and $\mathcal{S}$ is a safety risk functional capturing constraints such as avoidance of stimulation in uncertain artifact-dominated states or avoidance of parameters associated with adverse effects. The decision problem is partially observed, so policies depend on b_t. A risk-averse belief-space objective is

$$J(\pi) = \mathbb{E}_\pi\Big[\sum_{t\ge 1}\gamma^{t-1}\,\mathbb{E}_{(s_t,r_t,A_t)\sim b_t}[\ell(s_t, r_t, u_t)]\Big], \tag{4.3}$$

with discount γ. For safety, it is beneficial to penalize not only expected loss but also uncertainty [23]. We therefore incorporate a conservative uncertainty penalty, for example by adding a term proportional to the posterior variance of seizure risk or of action effects. If $\Delta(u_t)$ denotes the predicted reduction in seizure risk from action u_t under the model, then a conservative utility can be

$$\mathcal{U}(u_t \mid b_t) = \mathbb{E}_{b_t}[\Delta(u_t)] - \eta\sqrt{\mathrm{Var}_{b_t}(\Delta(u_t))} - \lambda_{\text{stim}}\|u_t\|^2 - \lambda_{\text{risk}}\mathbb{E}_{b_t}[\mathcal{S}], \tag{4.4}$$

and the policy selects $u_t = \arg\max \mathcal{U}(u_t \mid b_t)$ among feasible actions. This resembles an upper-confidence-bound principle in reverse, using lower confidence on benefit to reduce risk under uncertainty.

Selecting stimulation targets is an instance of action selection with structured actions. Suppose there is a set of candidate nodes $\mathcal{K}$ where stimulation can be applied, and an action corresponds to choosing $k \in \mathcal{K}$ and parameters α. The predicted effect of stimulating node k depends on controllability properties of the inferred directed graph and on the current regime. In a linearized approximation around the current belief mean $\mu_t = \mathbb{E}[s_t]$, the controlled dynamics are $s_{t+1} \approx (F_{r_t} + A_t)\,s_t + G_k\alpha + \xi_t$, where G_k is the actuation vector for node k. A first-order estimate of how stimulation changes expected seizure risk can be derived by differentiating a risk surrogate with respect to the mean state [24]. Let $\rho(s_t)$ be a differentiable surrogate for seizure propensity. Then the expected immediate change under action $u_t = (k, \alpha)$ is approximated by

$$\mathbb{E}[\rho(s_{t+1}) \mid b_t, u_t] - \mathbb{E}[\rho(s_{t+1}) \mid b_t, u_t = 0] \approx \nabla\rho(\bar{s}_{t+1})^\top G_k\alpha, \tag{4.5}$$

where $\bar{s}_{t+1}$ is the predicted next-state mean under no stimulation. This provides a computationally cheap scoring rule for candidates k when budgets are tight. Higher-budget configurations can use short rollouts to capture nonlinearities and regime transition effects.

For bilateral or multifocal seizures, the action selection must incorporate hemispheric structure [25]. If nodes are partitioned into left and right sets, one may compute directed influence imbalances under the belief state and prioritize stimulation on the side with higher inferred seizure-driving impact, while maintaining uncertainty penalties. Intracranial network analyses have proposed extracting directed connectivity differences between hemispheres and selecting stimulation targets from the more impactful side, with reported alignment with clinician decisions and associations between connectivity differences, seizure types, and stimulation outcomes [26]. In our framework, such hemisphere-dependent decision cues are treated as belief-state functionals of A_t and r_t, and they can be computed in high-resolution modes to inform long-term target planning, while online closed-loop actions can use simplified approximations consistent with compute budgets.

Safety constraints can be enforced through barrier-like penalties in belief space. Define a safety set $\mathcal{B}$ of beliefs for which action is considered safe, for example beliefs where artifact probability is low and where uncertainty is below a threshold. A soft barrier can be constructed using a function $\varphi(b_t)$

that grows large near the boundary of $\mathcal{B}$. The action utility can incorporate $-\lambda_{\text{bar}}\varphi(b_t, u_t)$, discouraging actions that push the predicted next belief outside $\mathcal{B}$. In practice, one approximates the next belief through a one-step prediction under the chosen filter approximation. This approach is compatible with compute constraints because it can use coarse uncertainty summaries such as regime entropy and artifact indicators rather than full belief distributions [27].

The online system therefore consists of three coupled components: a budgeted filter that updates b_t with adaptive gating, a forecaster that estimates $P_{t,H}$ under belief propagation, and a safety-aware controller that selects u_t by conservative utility maximization under feasibility and budget constraints. Because the same latent representation is used across modes, the system can incorporate high-resolution clinical information as priors and can maintain semantic consistency between wearable alerts and clinical intervention planning.

5. Theoretical Properties and Robustness Considerations

This section discusses analytical properties of the proposed approach, focusing on stability of compressed filtering, the role of invariance regularization in reducing spurious dependencies, and robustness under artifacts, missingness, and model mismatch. The aim is not to claim perfect guarantees in a complex physiological domain, but to articulate conditions under which the approach is expected to behave predictably and to identify failure modes that can be monitored.

A fundamental requirement for deployment is that the approximate filter remains stable under compression and gating. Consider a simplified linear-Gaussian regime with fixed $r_t = r$ over a time interval, and suppose the true dynamics are $s_{t+1} = Fs_t + \xi_t$ and the observation is $x_t = Hs_t + \epsilon_t$. The exact Kalman filter produces a covariance recursion $P_{t+1} = FP_tF^\top + \Sigma - FP_tH^\top(HP_tH^\top + R)^{-1}HP_tF^\top$. A compressed filter replaces P_t with $\tilde{P}_t$ in a restricted family, such as diagonal plus low-rank. The deviation $\Delta_t = \tilde{P}_t - P_t$ evolves under a nonlinear map influenced by the approximation projection. If the projection is non-expansive in an appropriate metric and if the pair (F, H) is detectable, then one can bound the growth of $\|\Delta_t\|$ by a function of the projection error and noise levels [28]. While deriving exact bounds depends on the projection choice, the key point is that stability is facilitated when approximation errors are uniformly bounded and when the filter update preserves positive semidefiniteness, which can be enforced by constructing $\tilde{P}_t$ as $D_t + L_tL_t^\top$ with $D_t \succeq 0$. In the budgeted setting, gating reduces observation information at some times. Stability then requires that informative updates occur sufficiently often, which is consistent with belief-driven gating that increases feature computation when uncertainty rises.

For switching regimes, stability depends on the ability to track r_t and to avoid mixture explosion. If the regime process is such that transitions are not excessively rapid relative to observation informativeness, then posterior regime probabilities can concentrate, allowing the filter to maintain a small number of active components. A practical theoretical surrogate is to require that the expected log-likelihood ratio between the true regime and alternatives is sufficiently large over typical windows, so that Bayesian updates contract uncertainty. Counterfactual consistency regularization can improve this by discouraging reliance on fragile modality-specific features that yield misleading likelihood ratios when sensors change [29].

We next consider the role of counterfactual consistency regularization in suppressing spurious dependencies. In electrophysiology, spurious dependencies can arise from volume conduction, shared artifacts, and montage-dependent mixing. If a predictor is trained on a single montage, it can learn to rely on features that correlate with seizures only because of that montage's mixing patterns. The invariance penalties $\mathcal{R}_{\text{cf}}$ and $\mathcal{R}_{\text{post}}$ act as constraints that favor representations whose risk predictions remain stable under perturbations that emulate montage changes and channel dropout. From a causal perspective, these penalties encourage the model to base predictions on latent variables that are more likely to correspond to stable physiological mechanisms rather than on measurement-specific proxies. While this does not prove causal identification, it can reduce distribution shift sensitivity, which is critical for real-world monitoring where electrode positions and impedances drift [30].

Artifacts are a dominant robustness challenge. Artifact processes can be heavy-tailed, bursty, and correlated across channels. The model includes explicit artifact terms $a_t^{(m)}$, but fully inferring them can be computationally expensive. A robust alternative is to use heavy-tailed likelihoods. For instance, modeling observation noise as Student-t corresponds to a Gaussian scale mixture, introducing per-time latent scales that downweight outliers. This can be integrated into variational filtering with modest cost [31]. Another approach is to infer an artifact probability indicator that gates channel trust. Both are compatible with budgeted inference because they add only a small number of scalar variables, and they provide a principled mechanism for the safety-aware controller to avoid acting when artifact probability is high.

Missing data and channel dropout are handled naturally by gating and by observation models that condition on available features. However, systematic missingness can still bias inference, for example if certain channels fail more often during motion that coincides with seizure events. Counterfactual perturbation training partially addresses this by exposing the model to channel dropout during training, but additional safeguards include monitoring shifts in the inferred noise covariance $R_t^{(m)}$ and increasing uncertainty when measurement quality degrades. In belief-space control, this increase in uncertainty reduces intervention aggressiveness through the conservative utility [32].

Model mismatch is inevitable, particularly in the directed graph component A_t. Directed interactions inferred from data may reflect effective connectivity under the model rather than direct causation, and the inferred directionality can be sensitive to preprocessing and to the choice of interaction features. The proposed approach mitigates this by treating A_t as a latent variable with structured priors and by focusing on decision-relevant functionals that average over posterior uncertainty. For example, rather than committing to a single best adjacency, target selection evaluates expected utility under the posterior over A_t, and uncertainty penalties discourage actions that rely on brittle directionality estimates.

The compute-constrained objective introduces a dual robustness issue: under severe budgets, the filter may become underpowered and lose the ability to track transitions. The method addresses this through adaptive gating, but it is still necessary to ensure that worst-case compute limits do not prevent any meaningful updates [33]. One can analyze a minimal informativeness condition: there should exist a baseline feature set of cost within $C_{\max}$ such that the mutual information between observations and regimes is non-negligible. This condition is application- and patient-dependent, and it motivates personalization. If a patient's seizures manifest strongly in a small number of channels or features, then a low budget may suffice, whereas other patients may require higher cost to reach acceptable calibration. The budgeted objective does not eliminate this reality, but it makes the trade-off explicit and measurable.

Finally, safety-aware control introduces a conservatism trade-off. Overly conservative uncertainty penalties can reduce intervention frequency and potentially reduce therapeutic benefit [34]. The appropriate conservatism level depends on clinical context and device goals, and it can be tuned using off-policy evaluation on recorded data or using controlled clinical studies. The benefit of the framework is that conservatism is parameterized explicitly through η and barrier penalties, and its effects can be analyzed systematically rather than being implicit in ad hoc thresholds.

6. Experimental Protocol, Metrics, and Deployment Constraints

This section outlines how the proposed framework can be evaluated and validated in a way that reflects its dual goals: accurate forecasting under resource constraints and decision-relevant modeling for intervention planning. The emphasis is on protocol design that separates training, personalization, and evaluation, and that reports trade-offs rather than single numbers.

For forecasting, the primary outputs are calibrated risk probabilities over a horizon and regime posterior trajectories [35]. Evaluation therefore should include both discrimination metrics and calibration metrics. Discrimination metrics include sensitivity at clinically relevant false alarm rates, time-to-detection relative to annotated seizure onset, and event-based precision-recall measures that account for seizure clustering. Calibration metrics include proper scoring rules such as negative log likelihood

of seizure events and Brier score on risk forecasts. Because the model produces uncertainty, one should also report how forecast confidence aligns with realized outcomes, for example by grouping predictions into quantiles and comparing predicted and observed event rates. These calibration analyses are especially important when the output is used to trigger stimulation or alerts, because overconfident false alarms can impose burdens and risks [36].

Resource constraints must be evaluated explicitly. The framework produces a family of models and inference configurations parameterized by compute budgets, gating aggressiveness, and graph capacity. Evaluation should therefore report performance as a function of measured compute, such as average feature computations per second, average latency per update, and energy proxies. One should also report worst-case compute and latency, because real-time devices must meet strict deadlines. Belief-driven gating can reduce average cost but may increase variance; therefore the evaluation should include tail metrics, such as the 99th percentile latency, to ensure gating does not violate real-time constraints during high-uncertainty episodes.

For cross-modal consistency, evaluation can be performed in settings where both modalities are available for at least some windows, or where one modality can be treated as a refinement reference [37]. One useful test is to fit the model in a low-resolution wearable configuration on scalp data and then assess whether the inferred coarse directed motifs align with motifs derived from intracranial refinement when aggregated to the same resolution. Alignment does not require exact equality, but it should be better than chance and stable across seizures. Another test is to evaluate how forecasts trained on one montage or session transfer to another montage or session. Counterfactual consistency regularization is intended to reduce the performance drop under such shifts, so the protocol should include deliberate montage perturbations and channel dropout scenarios.

For intervention planning, evaluation depends on available outcome data [38]. When stimulation outcome data are available, one can test whether model-derived target utilities correlate with seizure reduction under stimulation, while accounting for confounders and limited sample sizes. When only retrospective planning decisions are available, one can test agreement between high-utility target sets and clinical targets, but agreement alone is not definitive. More informative is stability: whether the top-ranked targets remain top-ranked across posterior samples and across different seizures within the same patient. Stability indicates that the model is not merely fitting transient patterns. The framework's uncertainty-aware scoring enables such analyses directly [39].

Deployment constraints include artifact prevalence, missing channels, and nonstationarity. Protocols should include stress tests that simulate motion artifacts, electrode detachments, and impedance drift, and should report how uncertainty estimates respond. A desirable behavior is that uncertainty increases and the safety-aware controller becomes more conservative when data quality degrades, rather than producing confident but unreliable forecasts. Additionally, one should test duty-cycled operation where the device deliberately reduces sampling or feature computation for power savings, and assess whether adaptive gating compensates effectively during high-risk periods.

Personalization protocols should distinguish between initialization, adaptation, and evaluation [40]. A realistic scenario is to initialize from population parameters, adapt patient-specific parameters using a limited calibration period, and then evaluate on held-out long-horizon data. Adaptation should be constrained to avoid drift, and evaluation should report both immediate post-adaptation performance and longer-term stability. Because seizure patterns can change over months, one should also evaluate incremental adaptation strategies that update slowly while maintaining safety.

Finally, one must consider interpretability and clinical usability. Directed graphs and latent states can be difficult to interpret. The framework addresses this by producing decision-relevant summaries such as node influence scores, hemisphere imbalance measures, and uncertainty intervals, but these outputs must be presented in a way that supports clinician reasoning [41]. Importantly, the framework does not require that clinicians accept the entire graph as ground truth; rather, it provides quantitative cues with uncertainty that can be integrated with other information. In deployment, the primary safety requirement is that the system communicates uncertainty and avoids brittle behavior, especially when acting through stimulation.

7. Clinical and Engineering Implications

A practical framework must map clean mathematical objects to messy clinical realities. This section discusses how the proposed approach interfaces with clinical workflows, device engineering constraints, and the broader question of how to build systems that remain reliable across heterogeneous patient populations and sensing conditions.

A key clinical reality is that monitoring and planning occur on different timescales [42]. Wearable monitoring must operate continuously for months, while intracranial recordings occur in a limited presurgical window. The proposed approach uses the same latent semantics across both timescales by representing seizures as regime transitions in a hybrid state-space model and by representing propagation patterns through a directed graph that can be refined when higher-resolution data are available. In practice, this means that wearable monitoring can produce stable summaries such as personalized risk baselines, circadian modulation estimates, and coarse regional influence patterns, while intracranial refinement can focus on pinpointing contact-level targets and validating directional motifs. The invariance regularization helps ensure these outputs are not contradictory merely due to modality differences.

From an engineering perspective, compute budgets are not abstract [43]. Feature computation can dominate energy cost, especially for interaction features that require cross-channel operations. The gating mechanism enables a principled form of conditional computation: expensive features are computed only when their expected information gain is high. This is analogous to event-triggered control, but applied to sensing and inference. In hardware, this can be implemented through low-power always-on processing for baseline features and a higher-power co-processor for burst computation. The budgeted objective provides a training-time mechanism to shape this behavior rather than relying on hand-tuned heuristics [44].

Artifact handling is particularly important in stimulation devices because stimulation itself can create artifacts. The model's explicit artifact terms and robust likelihoods provide a mechanism to detect when the observation model is violated and to increase uncertainty. The safety-aware controller can then reduce stimulation probability during artifact-dominated periods. This reduces the risk of positive feedback loops where stimulation artifacts trigger further stimulation. The framework also suggests that stimulation actions should be treated as interventions in a causal sense: they change the system dynamics and the observation distribution. The belief-space formulation naturally accounts for this by conditioning predictions on actions and by updating the transition model accordingly [45].

Another implication concerns how to interpret directed connectivity. Clinicians often reason in terms of seizure onset zones and propagation pathways, but directed connectivity estimates can vary with method and preprocessing. The proposed framework does not claim that inferred A_t is a direct measurement of anatomical connectivity; rather, it is a latent variable that captures predictive influence under the model, with uncertainty. This is aligned with decision-making: one can act on high-confidence influence patterns while refraining from acting on ambiguous ones. The invariance penalties further encourage that whatever influence patterns are inferred are stable under plausible sensor perturbations, which increases their practical trustworthiness [46].

Personalization raises questions about data quantity and labeling. Continuous wearable data are abundant but weakly labeled; intracranial data may be more precisely annotated but limited. The framework supports semi-supervised learning by leveraging the generative model and by using regime labels when available, without requiring exhaustive labeling. Additionally, personalization updates can be constrained to avoid overfitting to short-term idiosyncrasies. This is crucial because patients can exhibit rare seizure morphologies, and a model that overfits to recent patterns may miss those rare events [47]. Conservative adaptation and uncertainty-aware forecasting help mitigate this risk.

Finally, the framework suggests a pathway for integrating prediction with planning without forcing all decisions into a single monolithic model. The latent model provides a belief state; planning can use belief functionals and conservative utilities; clinicians can interpret summaries and uncertainties. This separation supports modular validation. The monitoring component can be validated for calibration

and stability under artifacts and shifts, while the planning component can be validated for robustness and alignment with outcomes [48]. Such modularity can be beneficial when deploying in safety-critical settings where incremental adoption is common.

8. Conclusion

This paper developed a causal-consistent, budgeted framework for cross-modal seizure forecasting and closed-loop intervention that unifies continuous monitoring and therapy-relevant modeling within a single latent state-space identification approach. The central technical idea is to learn and deploy a hybrid latent model with a sparse, time-varying directed interaction structure while explicitly optimizing under computational budgets and enforcing cross-modal stability through counterfactual consistency regularization. The resulting system supports adaptive feature and channel gating, uncertainty-calibrated forecasting, and belief-space, safety-aware action selection that is conservative under uncertainty and artifacts. The framework is designed to enable semantic continuity between low-cost wearable inference and high-resolution intracranial refinement, supporting personalization and clinically interpretable summaries without requiring that any one modality dominate the representation. By making compute and invariance constraints explicit in the learning objective and by tying intervention choices to posterior uncertainty, the approach provides a principled basis for building seizure technologies that remain robust under real-world sensing variability and that better align monitoring outputs with intervention planning needs [49].

References

[1] J. Clasadonte, T. Deprez, G. S. Stephens, G. Mairet-Coello, P.-Y. Cortin, M. Boutier, A. Frey, J. Chin, and M. Rajman, ''Δfosb is part of a homeostatic mechanism that protects the epileptic brain from further deterioration.,'' *Frontiers in molecular neuroscience*, vol. 16, pp. 1324922–, 1 2024.

[2] A. Vasquez, K. Islam, M. R. Cross, K. J. Miller, J. J. V. Gompel, B. N. Lundstrom, and A. L. Fine, ''Radiofrequency thermocoagulation in focal epilepsy: A retrospective cohort study.,'' *Epilepsia open*, vol. 10, pp. 529–538, 2 2025.

[3] S. S. Rabidas, C. Prakash, J. Tyagi, J. Suryavanshi, P. Kumar, J. Bhattacharya, and D. Sharma, ''A comprehensive review on anti-inflammatory response of flavonoids in experimentally-induced epileptic seizures.,'' *Brain sciences*, vol. 13, pp. 102–102, 1 2023.

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

[5] M. Charpentier-Hélary, A. de la Chapelle, M. Linard, N. André-Obadia, S. Boulogne, H. Catenoix, J. Jung, S. Rheims, K. Schiller, B. Frauscher, P. Ruby, and L. Peter-Derex, ''Dreaming in patients with epilepsy: a cross-sectional cohort study.,'' *Journal of sleep research*, vol. 34, pp. e14464–, 1 2025.

[6] P. Müller, D. S. Takacs, U. B. S. Hedrich, R. Coorg, L. Masters, K. E. Glinton, H. Dai, J. A. Cokley, J. J. Riviello, H. Lerche, and E. C. Cooper, ''Kcna1 gain-of-function epileptic encephalopathy treated with 4-aminopyridine.,'' *Annals of clinical and translational neurology*, vol. 10, pp. 656–663, 2 2023.

[7] S. Rodriguez, S. Sharma, G. Tiarks, Z. Peterson, K. Jackson, D. Thedens, A. Wong, D. Keffala-Gerhard, V. B. Mahajan, P. J. Ferguson, E. A. Newell, J. Glykys, T. Nickl-Jockschat, and A. G. Bassuk, ''Neuroprotective effects of naltrexone in a mouse model of post-traumatic epilepsy,'' 10 2023.

[8] G. Erdemir and A. N. Moosa, ''Electroclinical features of infantile epileptic spasms syndrome.,'' *Annals of Indian Academy of Neurology*, vol. 27, pp. 227–235, 5 2024.

[9] Y. Chen, B. Litt, F. Vitale, and H. Takano, ''On-demand seizures facilitate rapid screening of therapeutics for epilepsy.,'' 4 2025.

[10] O. S. Abdelaziz and A. A. F. D. Salles, *Hypothalamic Hamartoma-Related Epilepsy*, pp. 709–713. Springer International Publishing, 2 2023.

[11] E. Braun, F. M. Gualano, P. Siddarth, and E. Segal, ''Second-line cannabis therapy in patients with epilepsy.,'' *Clinical neurology and neurosurgery*, vol. 227, pp. 107638–107638, 2 2023.

[12] A. L. Muller, L. Diaz-Arias, M. C. Cervenka, and T. J. W. McDonald, ''The effect of anti-seizure medications on lipid values in adults with epilepsy.,'' *Epilepsy & behavior : E&B*, vol. 144, pp. 109260–109260, 5 2023.

[13] A. Leviton and T. Loddenkemper, ''Improving the health literacy of persons with epilepsy.,'' *Epilepsy & behavior : E&B*, vol. 163, pp. 110237–110237, 12 2024.

[14] A. Ghuli, A. Kannur, A. Mali, and A. Mangasuli, ''Quantification of eeg characteristics for epileptic seizure detection and monitoring of anaesthesia using spectral analysis,'' *International Journal of Intelligent Systems and Applications*, vol. 16, pp. 40–52, 4 2024.

[15] J. J. Pearce and R. W. Byrne, ''Book review: Posttraumatic epilepsy: Basic and clinical aspects,'' *Neurosurgery*, vol. 92, pp. e111–e112, 3 2023.

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

[17] C. Graham, M. Chladek, A. Nacson, L. Batchelder, L. Gleeson, A. Buenfil, E. Kim, O. Meyers, and B. Williams, ''Trends in coa use and regulatory acceptance for epileptic encephalopathy and developmental epileptic encephalopathy indications (p13-1.006),'' *Neurology*, vol. 100, 4 2023.

[18] M. K. Irwin, A. D. Patel, H. Palme', D. M. Cohen, C. Jones, and D. Skinner, ''Parental perceptions of school experiences for children with epilepsy.,'' *Journal of child neurology*, vol. 40, pp. 332–341, 1 2025.

[19] J. Thomas, C. Abdallah, Z. Cai, K. Jaber, J. Gotman, S. Beniczky, and B. Frauscher, ''Investigating current clinical opinions in stereoelectroencephalography-informed epilepsy surgery.,'' *Epilepsia*, vol. 65, pp. 2662–2672, 8 2024.

[20] A. Coppola, S. Krithika, M. Iacomino, D. Bobbili, S. Balestrini, I. Bagnasco, L. Bilo, D. Buti, S. Casellato, C. Cuccurullo, E. Ferlazzo, C. Leu, L. Giordano, G. Gobbi, L. Hernandez-Hernandez, N. Lench, H. Martins, S. Meletti, T. Messana, V. Nigro, M. Pinelli, T. Pippucci, R. Bellampalli, B. Salis, V. Sofia, P. Striano, S. Striano, L. Tassi, A. Vignoli, A. E. Vaudano, M. Viri, I. E. Scheffer, P. May, F. Zara, and S. M. Sisodiya, ''Dissecting genetics of spectrum of epilepsies with eyelid myoclonia by exome sequencing.,'' *Epilepsia*, vol. 65, pp. 779–791, 12 2023.

[21] J. Knowles, ''Attentional deficits and absence epilepsy: A tale of 2 interneuronopathies.,'' *Epilepsy currents*, vol. 24, pp. 188–190, 5 2024.

[22] A. Hyslop and M. Fajardo, ''Neuromodulation in pediatric drug-resistant epilepsy.,'' *Epilepsy & behavior : E&B*, vol. 165, pp. 110332–110332, 2 2025.

[23] M. R. Silverman, M. L. Smith, W. S. MacAllister, N. Heydari, R. M. Busch, R. Fee, and M. J. Hamberger, ''36 naming in monolingual and bilingual children with epilepsy,'' *Journal of the International Neuropsychological Society*, vol. 29, pp. 35–36, 12 2023.

[24] C. E. Hershberger, S. Louis, R. M. Busch, D. Vegh, I. Najm, P. Bazeley, C. Eng, L. Jehi, and D. M. Rotroff, ''Molecular subtypes of epilepsy associated with post-surgical seizure recurrence.,'' *Brain communications*, vol. 5, pp. fcad251–, 8 2023.

[25] G. Gupta, E. Saylam, A. Patel, L. Dang, L. O'Brien, S. Ibadat, J. Sturza, K. Neff, and R. Shellhaas, ''0773 behavioral sleep characteristics may differ in toddlers with epilepsy compared to children without epilepsy,'' *SLEEP*, vol. 46, pp. A341–A341, 5 2023.

[26] G. Peng, M. Nourani, and J. Harvey, ''Seeg-based bilateral seizure network analysis for neurostimulation treatment,'' *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2025.

[27] A. Taga, I. Cheong, K. E. Green, and M. D. Kornberg, ''Meningeal dissemination and drop metastasis from glioma presenting with non-epileptic myoclonus and minipolymyoclonus.,'' *The Neurohospitalist*, vol. 15, pp. 271–274, 10 2024.

[28] A. Min, W. R. T. Miller, K. Connelly, C. Nippert-Eng, H.-C. Lin, and P. C. Shih, ''Enhancing epilepsy awareness and cooperative care in elementary and middle schools,'' *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, pp. 1–25, 4 2024.

[29] V. R. Subramaniam, L. Mu, and C.-S. Kwon, ''Comparing vagus nerve stimulation and resective surgery outcomes in patients with co-occurring autism and epilepsy to patients with epilepsy alone: A population-based study.,'' *Autism research : official journal of the International Society for Autism Research*, vol. 16, pp. 1924–1933, 8 2023.

[30] F. Golab, P. Hajimirzaei, S. Zarbakhsh, S. Zolfaghari, P. Hayat, M. T. Joghataei, F. Bakhtiarzadeh, and N. Ahmadirad, ''Interplay of neuroinflammation and epilepsy in glioblastoma multiforme: Mechanisms and therapeutic implications.,'' *Journal of molecular neuroscience : MN*, vol. 75, pp. 68–, 5 2025.

[31] S. Magaki, M. Haeri, L. J. Szymanski, Z. Chen, R. Diaz, C. K. Williams, J. W. Chang, Y. Ao, K. L. Newell, N. Khanlou, W. H. Yong, A. Fallah, N. Salamon, T. Daniel, J. Cotter, D. Hawes, M. Sofroniew, and H. V. Vinters, ''Hyaline protoplasmic astrocytopathy in epilepsy.,'' *Neuropathology : official journal of the Japanese Society of Neuropathology*, vol. 43, pp. 441–456, 5 2023.

[32] G. Tamula, R. Katyal, S. Beniczky, and F. A. Nascimento, ''Focal cortical dysplasias: New advances for curing epilepsy.,'' *Epileptic disorders : international epilepsy journal with videotape*, vol. 25, pp. 284–284, 5 2023.

[33] S. N. Mbalinda, M. Kaddumukasa, J. N. Najjuma, D. R. Birungi, M. N. Kaddumukasa, J. B. Levin, C. H. Still, C. J. Burant, A. C. Modi, E. T. Katabira, and M. Sajatovic, ''Self-management intervention for reducing epilepsy burden among adult ugandans with epilepsy (smart-u): Randomised clinical trial protocol.,'' *Neuropsychiatric disease and treatment*, vol. 20, pp. 2277–2286, 11 2024.

[34] T.-H. Lee, Y.-C. Shih, Y.-J. Lu, C.-C. Chou, C.-C. Lee, H.-Y. Yu, and S.-J. Peng, ''Glucose metabolism of hippocampal subfields in medial temporal lobe epilepsy.,'' *Clinical nuclear medicine*, vol. 49, pp. 294–300, 2 2024.

[35] W. E. Kaufmann, ''Cdkl5 deficiency disorder: At the intersection between rett syndrome and developmental epileptic encephalopathies.,'' *Developmental medicine and child neurology*, vol. 66, pp. 410–411, 10 2023.

[36] E. L. Johnson, E. Bui, K. Tassiopoulos, M. O. Koretzky, R. Zepeda, E. Gonzalez-Giraldo, and R. F. Gottesman, ''Prevalence of epilepsy in people of sexual and gender minoritized groups.,'' *JAMA neurology*, vol. 81, pp. 996–996, 9 2024.

[37] R. M. Miralles, A. R. Boscia, S. Kittur, J. C. Hanflink, P. S. Panchal, M. S. Yorek, T. C. J. Deutsch, C. M. Reever, S. R. Vundela, E. R. Wengert, and M. K. Patel, ''Parvalbumin interneuron impairment causes synaptic transmission deficits and seizures in scn8a developmental and epileptic encephalopathy.,'' *JCI insight*, vol. 9, 10 2024.

[38] K. M. Kim and K. I. Yang, ''Sleep and epilepsy,'' *Sleep Medicine Research*, vol. 14, pp. 61–65, 6 2023.

[39] R. A. Sarkis, ''Epilepsy and amyloid: The plot thickens… so will the brain thin?,'' *Epilepsy currents*, vol. 25, pp. 266–268, 4 2025.

[40] D. M. Goldenholz, E. B. Goldenholz, and T. J. Kaptchuk, ''Quantifying and controlling the impact of regression to the mean on randomized controlled trials in epilepsy.,'' *Epilepsia*, vol. 64, pp. 2635–2643, 8 2023.

[41] G. Peng, M. Nourani, and J. Harvey, ''Seeg analysis to identify mis treatment for patients with focal epileptic seizures,'' in *2023 IEEE 23rd International Conference on Bioinformatics and Bioengineering (BIBE)*, pp. 187–194, IEEE, 12 2023.

[42] C. Lucasius, V. Grigorovsky, H. Nariai, A. S. Galanopoulou, J. Gursky, S. L. Moshe, and B. L. Bardakjian, ''Biomimetic deep learning networks with applications to epileptic spasms and seizure prediction.,'' *IEEE transactions on bio-medical engineering*, vol. 71, pp. 1056–1067, 2 2024.

[43] V. K. N. Shivamurthy, J. Gursky, and E. Yozawitz, ''A case of temporal lobe cavernoma causing epileptic spasms that resolved with surgical resection.,'' 10 2024.

[44] J. Y. Medina, J. Young, W. T. Miller, and A. Razi, ''Exploring online support needs of adolescents living with epilepsy,'' in *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing*, pp. 558–564, ACM, 11 2024.

[45] H. Chen, A. L. Numis, R. A. Shellhaas, J. R. Mytinger, D. Samanta, R. K. Singh, S. A. Hussain, D. Takacs, K. G. Knupp, L.-R. Shao, and C. E. Stafstrom, ''Treatment efficacy for infantile epileptic spasms syndrome in children with trisomy 21.,'' *Frontiers in pediatrics*, vol. 13, pp. 1498425–, 2 2025.

[46] J. Zhou, P. Guo, S. G. Karicheri, A. Sano, and Z. Haneef, ''Novel method to classify epileptic and non-epileptic seizures using wearable watch data (p5-9.007),'' *Neurology*, vol. 104, 4 2025.

[47] R. Garg, A. Pinard, and R. Wallerstein, ''A novel usp51 variant in a patient with autism spectrum disorder and epilepsy,'' *Egyptian Journal of Medical Human Genetics*, vol. 25, 11 2024.

[48] V. Jaramillo-Martinez, S. R. Sennoune, E. B. Tikhonova, A. L. Karamyshev, V. Ganapathy, and I. L. Urbatsch, ''Molecular phenotypes segregate missense mutations in slc13a5 epilepsy.,'' 5 2024.

[49] A. Aljishi, B. E. Sherman, D. M. Huberdeau, S. Obaid, A. Sivaraju, N. B. Turk-Browne, and E. C. Damisah, ''Statistical learning in epilepsy: Behavioral, anatomical, and causal mechanisms in the human brain.,'' 4 2023.