A Deployment-Oriented Review of EEG Signal Processing, Network Analysis, and Neuromodulation in Epilepsy

Zhuoran Xu[1], Sepehr Shirani[1,2], Saeid Sanei[3,4], Gonzalo Alarcón[5,6], Jiaxin Lei[7] , Antonio Valentín Huete[1,5*] and Ioannis Stavropoulos[1,5*]

[1]Department of Basic and Clinical Neuroscience, Institute of Psychiatry, Psychology and Neuroscience, King's College London

[2]Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Queen Square Institute of Neurology and Mental Health Neuroscience Department, Division of Psychiatry, University College London

[3]Department of Electrical and Electronic Engineering, Imperial College London

[4]Faculty of Brain, AI and Robotics, VinUniversity

[5]Department of Clinical Neurophysiology, King's College Hospital NHS Foundation Trust

[6]Department of Clinical Neurophysiology, Royal Manchester Children's Hospital

[7]School of Biomedical Engineering, Tsinghua University

[*] Shared last authorship

Abstract:

Approximately one third of patients with epilepsy remain refractory to pharmacological treatment. This creates a strong clinical need for methods that can support diagnosis, monitoring, and intervention under real clinical conditions. Despite substantial recent progress in epilepsy-related technologies and algorithms, strong retrospective performance often fails to translate into reliable chronic deployment because of drift, class imbalance, stimulation artifacts, and limited or absent ground-truth feedback.

This review provides a deployment-oriented synthesis of epilepsy-related methods across the sensing-to-control pipeline, spanning data acquisition, preprocessing, spontaneous event and evoked-response analysis, connectivity analysis, brain modeling, and closed-loop control. It is based on a structured scoping search, with studies selected and synthesized according to deployment relevance, validation setting, and role within the sensing-to-control pipeline. The methods are organized using a shared four-axis framework: temporal scale and latency constraints, observability and representation, vulnerability to non-stationarity and drift, and deployment role. This framework highlights how upstream constraints shape the reliability of downstream inference and intervention.

We argue that algorithm-guided neuromodulation in epilepsy is best understood as a constrained control problem under partial observability and non-stationarity. Within this framing, we revisit three recurring questions central to system design: whether seizure prediction is achievable, the validation status of candidate biomarkers, and the balance between focal and network models of epileptogenicity.

For each pipeline stage, we identify reporting and evaluation priorities aligned with clinically relevant endpoints rather than offline accuracy alone. Together, this review highlights engineering priorities for clinically durable systems and generates testable hypotheses about seizure dynamics and network controllability.

Key-words: Epilepsy; Neuromodulation; Closed-loop control; EEG Signal processing; Connectivity analysis; Brain modeling

1 Introduction

Epilepsy affects over 50 million people worldwide and remains refractory to pharmacological treatment in approximately one third of patients. This creates a strong clinical need for reliable methods to support diagnosis, monitoring, and treatment. In recent years, there has been rapid progress in technologies and algorithms for analyzing epileptic activity and supporting intervention. However, a substantial translational gap remains between strong offline performance and reliable clinical use.

This translational gap is not simply a consequence of patient-to-patient heterogeneity. It also reflects a broader mismatch between offline study conditions and real-world clinical deployment. In many cases, this mismatch does not arise from a single methodological step, but from interactions across stages, from data acquisition and preprocessing to downstream inference and intervention. Choices made early in the pipeline can constrain what later stages are able to detect, interpret, and act upon. Therefore, these methods need to be treated as parts of a deployment pipeline rather than as separable optimization problems.

Recent reviews have highlighted related translational challenges, including limited generalizability, weak prospective validation, and poor alignment between reported metrics and clinical burden.[1–4] However, these discussions have largely remained focused on specific tasks or individual stages of the broader pipeline. Much less attention has been paid to how such constraints propagate across stages. To address this gap, this review takes a cross-stage perspective on epilepsy-related methods, focusing on how constraints introduced at earlier stages influence the interpretation, reliability, and deployment of later-stage algorithms.

Rather than cataloging methods stage by stage in isolation, we organize epilepsy-related methods along a sensing-to-control framework under realistic clinical constraints (Fig. 1). To support comparison across heterogeneous approaches, we use a shared four-axis framework spanning temporal and latency constraints, observability and representation, vulnerability to non-stationarity and drift, and deployment role (Table 1). We begin by outlining the methodological scope of the review. We then review electrophysiological recording modalities and datasets, followed by preprocessing and artifact suppression as determinants of reliable observability. Next, we consider spontaneous epileptic events and stimulation-evoked responses as the primary sensing layer for downstream inference. Connectivity analysis and brain modeling are then discussed as intermediate representations for localization, interpretation, and control, before we turn to closed-loop neuromodulation and online adaptation.

2 Methods and conceptual scope

This review takes a conceptual, engineering-oriented approach rather than a systematic or meta-analytic one. The aim is not exhaustive coverage, but a structured synthesis of representative methodological classes across the epilepsy sensing-to-control pipeline. Throughout, we prioritize practical implementability, validation under realistic operating conditions, and reliability during continuous use under safety constraints and long-term non-stationarity. Therefore, we do not aim for Preferred Reporting Items for Systematic Reviews and Meta Analyses (PRISMA)-style coverage, a comprehensive catalog of stimulation hardware, or an exhaustive treatment of theoretical models without observable decision variables or clear implications for sensing, inference, or control.

2.1 Selection criteria

We conducted a structured scoping search across PubMed, IEEE Xplore, ScienceDirect, and Google Scholar. The search covered publications primarily from 1990 to 2026, while incorporating earlier foundational studies where needed. Keywords spanned three broad domains: recording modalities (electroencephalography (EEG), electrocorticography (ECoG), local field potentials (LFPs), and stereo-electroencephalography (sEEG)); signal processing and inference targets (artifact handling, event detection, seizure forecasting, evoked-response quantification, source localization, and connectivity analysis); and computational frameworks (machine learning, deep learning, dynamical modeling, and closed-loop control).

We distinguished between two broad categories of evidence. Primary evidence was used to support claims concerning robustness, reliability, evaluation validity, and clinical or device-level feasibility, and was prioritized when evaluation conditions reflected realistic use, including continuous or long-duration recordings, chronologically separated validation, explicit latency or safety constraints, stimulation-while-sensing settings, or robustness analyses with direct implications for interpretation and device operation. Contextual evidence, including reviews, foundational methodological studies, and mechanistic work, was used to define major method families, clarify terminology, and provide background where deployment-aligned validation remains limited.

2.2 The shared four-axis framework

Throughout the review, methods are compared using a shared four-axis framework:

- Temporal scale and latency constraints: the time scale at which information is extracted or acted upon, ranging from milliseconds to days, together with whether causal, low-latency operation is required.
- Observability and representation: whether the relevant variables are defined at the sensor level, source level, as evoked-response summaries, or as latent states derived from dynamical models, and the inferential constraints associated with each representation.
- Vulnerability to non-stationarity and drift: susceptibility to short-term state dependence and long-term evolution during chronic recording and stimulation.
- Deployment role: whether a method is used for offline mapping or planning, continuous monitoring, prospective prediction of seizure risk, or real-time closed-loop control.

Together, these four axes provide a common basis for comparing the methods reviewed in the sections that follow in relation to shared operational constraints, expected evidence, and deployment role.

3 Data Acquisition and Datasets

Electrophysiological data define the empirical conditions under which methods can be developed, evaluated, and interpreted. Recording modality, spatial sampling, temporal coverage, annotation quality, and dataset structure jointly define the conditions under which methods are developed, validated, and compared. However, in epilepsy research, data are not only heterogeneous across studies in recording conditions, label quality, and temporal structure, but also unevenly available across tasks and deployment contexts. Large-scale retrospective scalp EEG (scEEG) datasets are far more readily available than those derived from chronic implantable devices or from stimulation-while-sensing studies. As a result, benchmarking often remains shaped more by data availability than by clinical deployment needs.

These properties determine which time scales are observable, which representations can be supported, which forms of non-stationarity and drift become visible during evaluation, and which deployment roles can be assessed with credibility. We first review the main recording modalities and their acquisition constraints. We then discuss representative public resources in terms of the evidence they support, before considering how dataset properties propagate into downstream methodological bias and evaluation failure modes.

3.1 Recording modalities and acquisition constraints

The principal electrophysiological recording modalities in epilepsy are scEEG and intracranial EEG (iEEG), the latter encompassing ECoG and sEEG. These modalities differ not only in invasiveness, spatial resolution and coverage, signal-to-noise ratio (SNR), and feasible recording duration, but also in the kinds of inference they can support. In deployment terms, they provide complementary rather than interchangeable forms of evidence.

Intracranial EEG offers high spatial specificity and access to cortical and subcortical structures central to seizure onset localization, connectivity analysis, and stimulation studies. At the same time, it is shaped by patient-specific electrode placement, incomplete spatial coverage, and finite monitoring windows, all of which constrain comparability across subjects and limit inference beyond sampled regions.[5] In chronic neuromodulation settings, long-term non-stationarity is not merely a modeling nuisance but an intrinsic feature of the recording context.

Signals can evolve over months after implantation because of tissue response, impedance change, and physiological adaptation.[6] For these reasons, iEEG is often the most relevant sensing substrate for therapeutic neuromodulation, but it remains difficult to standardize across patients and cannot provide complete access to the network it seeks to characterize or control.

By comparison, scEEG provides broader spatial coverage and is more amenable to largescale data collection because it is non-invasive. This has made it the dominant substrate for population-scale datasets and many benchmarking efforts. However, broader coverage comes at the cost of lower spatial specificity, greater vulnerability to environmental and physiological artifacts, and weaker access to deep or highly focal generators. Therefore, its principal strengths lie in monitoring, screening, and large-scale model development, whereas its limitations become more apparent in settings that require precise localization, detailed connectivity analysis, or state estimation for stimulation.

Beyond recording modality, long-term and home-based EEG define a distinct acquisition context rather than simply extending standard scalp or intracranial recordings over longer periods. Recordings spanning months to years expose pronounced non-stationarity, incomplete event reporting, variable adherence, and human factors that affect annotation reliability and data completeness.[7–9] These properties matter directly for forecasting, calibration, and adaptive neuromodulation because they determine which failure modes become visible during evaluation. Therefore, chronic ambulatory and implantable-device recordings are especially important for assessing long-term forecasting realism and adaptive control. Stimulation-while-sensing data are similarly essential for single-pulse electrical stimulation (SPES), evoked-response quantification, and perturbational effective connectivity. In both cases, such datasets remain scarce.

3.2 Public datasets, shared resources, and standardization

Public datasets in epilepsy differ not only in size, but also in the kinds of methodological claims they can support. Some resources offer scale and heterogeneity but little chronic realism.

Others preserve longer temporal structure but remain small or difficult to access. Still others are better suited to more specific but clinically relevant targets.

Large clinical scEEG datasets remain among the most widely used public resources in epilepsy. The Temple University Hospital (TUH) EEG Corpus and its seizure-annotated subsets have become standard benchmarks because they expose algorithms to substantial variability in recording conditions and patient state.[10;11] Open infrastructures such as PhysioNet have also shaped the field by hosting canonical datasets, including the CHB-MIT Scalp EEG Database, and by supporting standardized formats, open tools, and reusable benchmarking practices.[12] More recently, structured clinical EEG resources such as the Harvard Electroencephalography Database (HEEDB) have added richer clinical metadata to this landscape.[13] These resources are well suited to benchmarking and large-scale retrospective analysis, but they cannot support stronger deployment-facing claims when label granularity, temporal continuity, signal quality, or clinical context are insufficient.

Longer-duration resources are especially important for seizure prediction because they preserve the temporal context needed to study pre-ictal structure, background dynamics, and time-aware validation.[14;15] EPILEPSIAE is a prominent example of this type of resource and NeuroVista remains notable as one of the few examples of truly chronic invasive recording used to study seizure forecasting over months to years.[14;16] Such datasets provide conditions under which circadian and multidien structure, calibration decay, and long-term non-stationarity can be examined in a meaningful way.[6;16;17] Their scarcity nevertheless limits how confidently long-term forecasting robustness can be assessed.[15]

Other public resources are organized around more specific targets. Some are designed for analysis of interictal epileptiform discharges (IEDs), including spike morphology, spatial field structure, and multi-channel consistency. A notable recent example is the dataset introduced by Lin et al., particularly given the relative scarcity of public EEG datasets with expert annotations and spatial information.[18] Others support sensing-while-stimulating paradigms, where stimulation-evoked responses are used to localize seizure onset zones and to characterize network behavior under controlled perturbation. In epilepsy, SPES-based datasets are one of the clearer examples of this kind of resource, and the PRIOS dataset provides a public example that combines intracranial recordings with explicit stimulation timing and protocol structure.[19] However, such datasets remain relatively scarce, especially when stimulation parameters, electrophysiology, and longitudinal clinical context are jointly represented.

Given the diversity of epilepsy datasets, standardized data organization and sharing frameworks are essential for meaningful comparison, reproducibility, and reuse across studies. Platforms such as iEEG.org facilitate multisite sharing of intracranial recordings together with imaging and clinical metadata, while also making visible the practical constraints imposed by governance, de-identification, and privacy.[20] The iEEG extension of the Brain Imaging Data Structure (iEEG-BIDS) provides a formal structure for organizing intracranial electrophysiology data and supports interoperability, scalable preprocessing, and cross-study reuse.[21]

3.3 Engineering consequences: how dataset properties bias methods and evaluation

From an engineering perspective, dataset properties impose structured constraints on method design and evaluation. Recording modality, spatial sampling, temporal continuity, and acquisition context shape the burden of artifact handling, the feasibility of continuous rather than clip-based modeling, and the need for subject-specific adaptation. They also influence whether calibration and threshold selection remain stable across sessions or devices, and which sources of drift or deployment-relevant failure modes can become visible during validation. Table 2 summarizes common dataset properties and their downstream methodological consequences across neuromodulation pipelines.

Annotation structure and granularity further determine which tasks can be formulated and how results should be evaluated. Event-level annotations may support continuous detection or temporal modeling, whereas segment-level labels are more naturally suited to clip-based classification. Channel-level, spatial, or state-dependent annotations further determine whether localization, propagation analysis, or context-aware modeling can be assessed. When these properties are underspecified, offline results become difficult to compare and easy to overinterpret, limiting their value for clinically durable neuromodulation systems.

4 Preprocessing and Artifact Suppression

Signal integrity defines the observability boundary for signal-processing and artificial intelligence (AI) methods used in epilepsy-related diagnosis, monitoring, and intervention. Neural signals can be contaminated by movement-, physiological-, and hardware-related artifacts and, in stimulation-based paradigms, by high-amplitude stimulation-locked artifacts. Therefore, preprocessing is not merely a technical prelude. It determines which features remain interpretable, which temporal windows remain observable, and which deployment-facing claims can be supported.

The discussion first considers general preprocessing across scEEG and iEEG, including filtering, referencing, quality control, and removal of conventional non-stimulation artifacts. It then turns to stimulation-artifact suppression when sensing and stimulation occur in the same system, with emphasis on method selection and on the interpretive limits imposed by early-window recovery.

4.1 General preprocessing across scEEG and iEEG

Although specific workflows differ across modalities and tasks, widely adopted preprocessing operations include filtering and referencing, channel and segment quality control, and removal of conventional non-stimulation artifacts.[22;23] The following discussion considers these components in turn, with emphasis on their implications for downstream analysis.

4.1.1 Filtering and referencing

Bandpass and notch filtering are routinely used to suppress slow drifts and power-line interference. However, in epilepsy and neuromodulation applications, filter choice is rarely neutral. For evoked-response analysis, spike detection, and seizure dynamics, phase distortion, transient smearing, and group delay can alter precisely those features on which downstream inference depends.[23] Even modest choices, such as a highpass cutoff of 0.1 versus 1 Hz, can reshape spike morphology, shift feature distributions, and modify connectivity estimates.

Referencing is similarly consequential. It affects spatial contrast, apparent synchrony, and the extent to which common-source contamination is preserved or attenuated. In scEEG, common-average and Laplacian montages are widely used under low SNR conditions.[22] In ECoG and sEEG, bipolar or other local references are often preferred to reduce spatial mixing and hardware-related contamination.[24] For cross-subject modeling and comparative analysis, robustness across plausible montages is often more informative than a single nominally optimal reference, because montage choice can materially alter biomarkers and network summaries.

4.1.2 Channel and segment quality control

Quality control aims to identify bad channels and corrupted segments before they distort downstream estimation. Common criteria include abnormal amplitude, variance, kurtosis, and line-noise dominance.[22] Iterative rejection and re-referencing can improve stability.[24] In iEEG, quality control must also account for modality-specific failure modes, including stimulation-induced saturation, seizure-related clipping, impedance fluctuation, and device-specific artifacts. These are not minor implementation details, since undetected failures can propagate into false connectivity structure, unstable detector thresholds, or spurious model drifts.

4.1.3 Conventional non-stimulation artifact removal

Ocular, muscle, cardiac, and slow-drift artifacts are commonly addressed using regression, time-frequency methods, blind source separation (BSS), and increasingly learning-based denoising approaches.[22–24] Independent component analysis (ICA) used for BSS remains widely used, including for component rejection in epileptic EEG and supervised denoisers have been trained to approximate ICA-cleaned signals.[25;26] These methods can improve signal quality in conventional settings, but they do not solve the more specific problems posed by stimulation-locked artifacts.

4.2 Stimulation-artifact suppression

Stimulation-induced artifacts differ qualitatively from conventional EEG contamination. They are pulse-locked, frequently exceed amplifier dynamic range, and often exhibit multiphasic waveforms shaped by electrode-tissue interface dynamics. Their morphology depends on pulse amplitude, pulse width, polarity, electrode geometry, impedance, and local tissue conditions, and they can obscure neural activity in the first milliseconds after stimulation. This is precisely the window required for quantifying early cortico-cortical evoked potentials (CCEPs) elicited by single-pulse electrical stimulation (SPES) and for estimating perturbational effective connectivity. Similar observability constraints also arise in deep brain stimulation (DBS) sensing and other invasive stimulation settings. Related though not identical problems also appear in transcranial electrical stimulation (tES), especially transcranial alternating current stimulation (tACS), and in transcranial magnetic stimulation electroencephalography (TMS-EEG).

Hardware-level mitigation can reduce artifact burden at acquisition, but it does not by itself guarantee recovery of the earliest usable post-stimulus signal. Its effectiveness remains constrained by safety, engineering complexity, and protocol dependence.[27;28] As a result, reliable recovery of early evoked activity often still depends on software-based artifact suppression. These methods can be grouped into a small number of broad families with distinct assumptions, operating characteristics, and failure modes (Table 3).

Conservative causal methods. Blanking, interpolation, and simple template subtraction remain attractive because they are straightforward, often causal, and robust under severe saturation or clipping. Their limitation is equally clear; they can remove or distort the early components required for latency-sensitive analysis. Even brief blanking windows may bias latency estimates, early amplitude measures, and connectivity-related conclusions in stimulation-EEG settings.[29] Template methods can be effective when artifacts are sufficiently stereotyped, but they lose reliability when waveform shape drifts across pulses, sessions, or impedance states.

Reconstruction, subspace, and model-based methods. When early-window fidelity matters, more elaborate approaches may be justified. Biophysical reconstruction has been used to recover early SPES-evoked responses and sparse approaches such as matching pursuit-based artifact reconstruction and removal method (MPARRM) attempt to separate artifact, neural signal, and noise with high temporal fidelity.[30;31] Projection- and subspace-based methods, including null-space projection, ICA, signal-space projection, and related variants, attempt to remove artifact-dominated components or spatial modes.[29;32] Parametric and state-space formulations model non-stationary artifact dynamics explicitly and may improve robustness under protocol variation.[33;34] Across these families, the main risks are over-removal of neural activity patterns, dependence on channel count or artifact coherence, and reduced robustness under protocol drift or model mis-specification.

Learning-based methods. Learning-based suppression methods can capture nonlinear artifact structure and may adapt across stimulation protocols.[34] However, they also introduce distinct risks, including over-smoothing, hallucinated structure, leakage between training and evaluation conditions, and failure of cross-site generalization.

4.2.1 Selection of stimulation-artifact suppression methods

Method selection should be driven by which properties of the post-stimulus signal must remain interpretable for the downstream endpoint, especially whether the analysis depends on the earliest post-stimulus activity or primarily on later components. If early post-stimulus activity is itself the primary endpoint, as in early CCEP quantification or perturbational effective connectivity, approaches that structurally remove or distort the early window should be avoided, because the relevant requirement is preservation of latency, sharp transients, and local waveform structure rather than nominal artifact-reduction performance. By contrast, when later components dominate the analysis, more conservative suppression strategies may be acceptable, provided that the discarded interval does not overlap with the features being measured.

Beyond task requirements, deployment context further constrains which suppression strategies are actually suitable. For real-time or closed-loop systems, causal and low-latency methods should be prioritized, and offline benchmarking pipelines should be clearly distinguished from pipelines that remain implementable under practical latency, power, and hardware constraints. For longitudinal or cross-session use, methods that rely on stable artifact morphology or highly tuned model assumptions may become less reliable as impedance, montage, stimulation parameters, and recording conditions drift over time. In that setting, robustness to non-stationarity may be more important than nominal suppression performance under controlled offline conditions.

4.2.2 Audit and claim boundaries for early-window interpretation

When early-response latency, early-response amplitude, or perturbational effective connectivity is interpreted from the first milliseconds after stimulation, artifact handling should be reported as an auditable interface rather than an implementation detail, because early window fidelity is a prerequisite for interpretation rather than an optional supplementary consideration. At minimum, studies should report blanking duration and any interpolation applied, state whether filtering was causal or acausal and whether group delay affected the analyzed early window, document whether saturation or clipping occurred and how quickly recordings recovered, quantify residual stimulus-locked artifact within the early window of interest, and describe how recovery quality varied with stimulation parameters or protocol conditions.

If early-window recovery is not demonstrated, or if the analyzed interval is structurally removed or heavily reconstructed, then strong claims about early-response biomarkers or perturbational effective connectivity should be qualified explicitly. Under those conditions, it is more defensible to restrict interpretation to later components or to frame the findings as protocol-conditioned response patterns rather than reliable evidence of early-pathway connectivity. More generally, reporting should make clear which aspects of the reported signal depend on deployment-compatible preprocessing and which depend on offline reconstruction or benchmarking-only procedures, so that the evidential boundary of the findings remains explicit.

5 Epileptic Event Characterization and Evoked Responses

Spontaneous epileptic activity and stimulation-evoked responses form the principal electrophysiological sensing signals for downstream epilepsy-related inference and intervention. These signals are not interchangeable. Spontaneous events index intrinsic excitability and pathological dynamics and provide time-stamped or state-like information for monitoring, triggering, and risk estimation. By contrast, stimulation-evoked responses probe tissue and network reactivity under controlled perturbation and provide structured summaries that can inform mapping, programming, and downstream control. The key issue at this stage is not simply whether a signal can be detected offline, but whether it can be extracted with sufficient reliability, temporal fidelity, and calibration to support downstream decisions under chronic non-stationarity.

We first consider spontaneous epileptic events, including IED detection, seizure detection, and seizure prediction, which place different demands on timing, representation, and deployment-oriented evaluation (Table 4). We then turn to stimulation-evoked neural responses, with a primary focus on SPES and its use in epilepsy-related sensing and inference.

5.1 Spontaneous epileptic events

Spontaneous epileptic activity comprises multiple abnormal phenomena. Here we focus on IED detection, seizure detection, and seizure prediction, because these represent three principal deployment-oriented sensing tasks with distinct outputs, time scales, and evidential requirements. IED detection is an event-timing problem at the millisecond scale; seizure detection is a state-transition problem on multi-second horizons; seizure prediction is a long horizon risk estimation problem operating over hours to days.

5.1.1 Task structure and methodological constraints

IED detection is constrained by morphological heterogeneity, low signal-to-noise ratio in scEEG, interrater disagreement, and cross-site or cross-device shift.[35] Concurrent scalp-intracranial recordings further show that some intracranial IED signatures remain only partially observable on scEEG, emphasizing that detectability is limited not only by model capacity but also by acquisition geometry.[36]

Seizure detection operates on longer time scales and must remain reliable under behavioral variability, changing vigilance state, and months-long non-stationarity.[37] Seizure prediction extends the problem further, aiming to estimate future seizure likelihood over minutes to days under extreme class imbalance, long-term drift, and inconsistent definitions of the pre-ictal state.[38–42] Ultra-long-term ambulatory and subcutaneous EEG studies make clear that these constraints are fundamental rather than incidental: forecasting must operate under sparse seizure occurrence, incomplete reporting, and months-to-years of biological and technical variation.[43–45]

Across these tasks, method families include rule- or morphology-based detectors and template-matching approaches that remain especially relevant to IED detection, conventional feature-based and machine-learning pipelines widely used in seizure detection and forecasting, and increasingly deep learning architectures that exploit temporal context and multichannel structure.[46–52] However, the broader methodological overlap does not remove the fact that the three tasks differ in output structure, time horizon, and deployment burden.

5.1.2 Evaluation and deployment implications

For IED detection, the relevant output is a time-stamped event, so low latency, onset jitter, and false positives per minute (FP/min) matter more than retrospective discrimination alone. In deployment, the main burden is the false positive rate during long continuous recordings, because clinician review scales directly with false-alarm frequency. Accordingly, the relevant evidence comes from long continuous streams with an explicit false-alarm budget and a stable operating point.

For seizure detection, the central trade-off is detection latency versus false alarms over long continuous streams. The main deployment challenge is maintaining that trade-off as background state, medication, vigilance, and signal characteristics change over time. Continuous stream evaluation at a stable operating point is therefore more informative than strong performance on isolated segments.

For seizure prediction, the relevant output is not a discrete event but a risk estimate or warning state. Therefore, the central endpoints include time in warning, false-warning burden, calibration, prediction timing, continuous-output behavior, and prospective leakage-aware validation.[53] Circadian and multidien structure may be among the more stable predictive signals in long-term settings, which is one reason probabilistic risk estimation can be more clinically plausible than brittle binary prediction.[16;17;40;54–57] The principal methodological risk is overestimating prospective utility from retrospective structure, particularly when temporal leakage, post-hoc pre-ictal windows, or unstable definitions of seizure risk inflate apparent predictability.[40;42;55;58] Therefore, the corresponding evidential standard is prospective validity, with calibration, warning burden, and strict temporal separation between training and evaluation treated as central rather than optional.[59–61]

Accordingly, deployment-relevant reporting should at minimum make explicit the label policy, temporal split logic, leakage or embargo controls, operating-point selection strategy, and continuous-stream error burden. Without these elements, favorable retrospective metrics remain difficult to interpret and weakly comparable across studies.

5.2 SPES-evoked neural responses

Among paradigms used to study stimulation-evoked responses in epilepsy, SPES is one of the most established frameworks for probing tissue and network reactivity under controlled perturbation. The resulting time-locked responses recorded at other sites can be used to probe functional organization, epileptogenicity, and perturbational effective connectivity.[62–66] Prospective studies have linked abnormal SPES-evoked responses to epileptogenic or structurally abnormal cortex and, in some settings, to postsurgical outcome, with similar localization value reported in pediatric cohorts.[67;68] Here, these responses are treated primarily as signal-level sensing objects; their translation into network-level representations is addressed in Section 6.

5.2.1 Definitions and component structure of SPES-evoked responses

In the literature, closely related terms such as SPES-evoked responses, cortico-cortical evoked potentials (CCEPs), and early or delayed responses (ER/DR) are not always used consistently. Here, SPES-evoked responses/CCEPs are used as the general term, while ER and DR refer to early and delayed components or response classes within them. Within SPES-evoked responses/CCEPs, ER and DR are best understood as two broad classes of evoked activity rather than perfectly discrete categories. ER generally refers to the earlier, relatively fixed-latency component, often associated with more direct pathways. DR refers to later and more variable responses, which may reflect broader network reverberation and can resemble spontaneous IED-like activity in epileptogenic tissue.[69–71] However, in practice, these distinctions are not always clear-cut. Conservative operational definitions of early and delayed responses are therefore important. Trial-level variability further motivates probabilistic or stability-aware definitions of response presence rather than binary response classification.[72;73] Methodological reviews now provide increasingly detailed guidance for defining and reporting ER/DR structure.[64;66;74]

5.2.2 Quantification and response representations

Time-domain summaries, including peak amplitude, peak-to-peak measures, latency, and area under the response curve, remain the most widely used representations and support constructions of stimulus-response matrices for mapping.[62;63;66] More structured representations have also been proposed. Basis profile curves (BPCs) and canonical response parameterization (CRP).[75;76] for example, aim to capture response shape using combinations of canonical waveform components. Transfer-function and response-shape approaches extend this further when evoked responses are used to construct downstream network objects.[77;78] Because SPES parameters themselves modulate morphology and recruitment, cross-session comparison is strongest when response summaries are interpreted together with protocol-aware normalization or explicit stimulation metadata.[79] These summaries are most defensible when latency, amplitude, and response shape remain stable enough for the intended use, especially when they are used to distinguish ER from DR or to construct downstream connectivity measures. Machine-learning approaches may support localization, response typing, and automated response detection or standardization.[78;80–83] However, they do not remove the requirement that the analyzed response window itself be sufficiently reliable for interpretation.

6 Connectivity Analysis and Networks

Epilepsy is increasingly treated as a network disorder. In this view, stimulation acts on nodes and pathways, while sensing and intervention depend on how activity propagates through patient-specific networks. Therefore, connectivity is clinically important in epilepsy-related inference and intervention. However, connectivity estimates remain especially vulnerable to misinterpretation because of spatial mixing and zero-lag confounds, dependence on preprocessing and referencing choices, and non-stationarity across vigilance, sleep, peri-ictal state, and longer-term drift. These constraints place strong limits on how connectivity measures can be interpreted and used. In this review, the central question is not connectivity estimation in the abstract, but which connectivity estimates remain robust enough to inform target ranking, stimulation programming, and outcome stratification.

We first distinguish SC, FC, and EC and define their interpretive boundaries. We then consider graph representations and dynamic network summaries, before turning to connectivity-informed target selection, stimulation programming, and the role of effective connectivity in decision support (Table 5).

6.1 Structural, functional, and effective connectivity; interpretability and robustness boundaries

SC, FC, and EC answer different questions and should not be treated as interchangeable. FC describes statistical dependence, such as correlation, coherence, or phase synchrony, without requiring an explicit causal or biophysical model. EC attempts to characterize directed influence under explicit perturbational or modeling assumptions.[84] SC provides anatomical constraints on which interactions are plausible, but it does not by itself specify when, how strongly, or in which direction those interactions will be expressed.

One important boundary on interpretability in epilepsy is state dependence. Vigilance, sleep stage, medication, and peri-ictal dynamics can alter coupling strength, directionality, and apparent topology even when anatomy is unchanged.[85] SC constrains propagation and stimulation spread, but it does not directly validate FC or EC, because expressed dynamics depend on synaptic gain, neuromodulatory state, and pathological synchronization.[86] Therefore, anatomical plausibility is an important constraint, but not a ground-truth oracle for electrophysiological connectivity.[85;86]

A further boundary is methodological rather than biological. Zero-lag correlations may arise from volume conduction or source leakage rather than true interaction. Leakage-aware measures such as imaginary coherence and weighted phase lag index (wPLI) reduce instantaneous confounds, but they may also down-weight genuine near-zero-lag physiology. [87;88] Source-space analysis does not eliminate this problem. Instead, it relocates it to inverse modeling and leakage correction, both of which can materially alter topology and reliability.[89] Referencing adds a further layer of sensitivity. Common-average, bipolar, and Laplacian montages can change edge weights, graph structure, and apparent hubs.[90] In iEEG, these issues are compounded by limited coverage and sampling geometry, which can create apparent hubs in densely sampled regions. Therefore, transparent reporting of coverage, distance-control analyses, and montage robustness is particularly important.

For deployment-facing use, interpretability alone is not sufficient. Connectivity estimates intended to support neuromodulation decisions should also demonstrate robustness to plausible referencing schemes, leakage-control choices, thresholding or graph-density variation, and physiological state.

6.2 Graph representations and dynamic network summaries

Graph representations provide a compact way to summarize connectivity in terms of node importance, segregation, integration, and pathway organization.[91] In epilepsy, hub-centered and network-disorder views have been influential because they align naturally with distributed models of seizure generation and propagation.[92] However, graph measures inherit the assumptions of the underlying connectivity estimator, together with the effects of referencing, parcellation, coverage, and thresholding. Therefore, they should be interpreted as summaries of a modeling pipeline rather than as direct readouts of network ground truth.

Graph-derived biomarkers become clinically useful only when they support stable decisions, such as ranking targets, stratifying patients, or improving outcome prediction.[93] Index-based summaries such as the connectivity epileptogenicity index (cEI) can provide decision-friendly orderings, but these orderings may still shift with preprocessing choices rather than stable pathophysiology.[94] Directed graphs derived from perturbational or propagation-based networks offer a more actionable representation and have supported epileptogenic zone localization in selected settings.[78] However, their value depends on robustness checks rather than on apparent graph coherence alone.

Static FC may fail to capture clinically relevant network dynamics because coupling varies with vigilance, sleep, medication, peri-ictal dynamics, and longer chronobiological structure.[85] Therefore, dynamic connectivity methods have been developed to model this non-stationarity through recurring network states, time-varying coupling, and transition structure.[95;96] Switching state-space models, including hidden Markov models (HMMs), are particularly attractive because they infer latent connectivity states without relying entirely on fixed sliding windows and can yield occupancy and transition summaries potentially useful for monitoring.[97] However, these summaries become clinically meaningful only if inferred states remain robust out of sample, correspond to interpretable physiological variation, and are not driven primarily by preprocessing choices or model assumptions.

6.3 Connectivity-informed target selection, stimulation programming, and effective connectivity

Network-guided neuromodulation reframes targeting as the attempt to steer pathological dynamics through patient-specific connectivity structure.[98;99] In practice, this requires combining different connectivity classes rather than privileging any single one. Structural information provides feasibility and safety constraints. Electrophysiological FC and EC can indicate which pathways are engaged under clinically relevant states. Directed measures, including approaches that approximate Granger causality, provide one route to directional inference. Perturbational and propagation-based measures can add directional information closer to programming decisions than undirected association alone. Accordingly, the most useful connectivity-guided strategies are those that integrate anatomical constraints, directional influence estimates, and state-dependent network information into a patient-specific decision framework.

The strongest translational evidence remains patient-specific rather than population-averaged. For example, in responsive neurostimulation, patient-specific connectivity has been linked to stimulation response more consistently than generic group-level network summaries.[100] Group-average networks remain useful for hypothesis generation and syndrome-level context. However, programming-relevant conclusions usually require within-patient validation because outcome-relevant pathways, stimulation spread, and stimulation-induced reconfiguration are highly individual.[101–104]

Perturbational EC provides protocol-conditioned directed edges that can be ranked or incorporated into stimulation programming constraints. Its main attraction lies in directionality grounded in controlled perturbation rather than inferred solely from passive covariance structure. However, edge existence and strength remain sensitive to stimulation parameters, preprocessing, and response definitions. Therefore, perturbational EC should be treated as a protocol-conditioned network object whose usefulness depends on within-session reproducibility and robustness to plausible preprocessing choices.

Directed pathways may also be inferred from reproducible temporal precedence during spontaneous spike or seizure spread. These propagation-derived summaries have been linked to surgical outcome and reinforce the idea that epileptogenic networks are defined not only by onset zones but also by recruitment structure.[105] However, their main limitations are detection bias, incomplete sampling, and marked state dependence. In practice, they are most informative when interpreted alongside perturbational EC and structural constraints, especially when apparent causality may in fact be bidirectional.

Taken together, connectivity is best viewed not as a ground-truth layer sitting above signal analysis, but as an intermediate representation whose value depends on stability across preprocessing choices, physiological states, and sampling limitations.

7 Brain Modelling and Cross-Modal Transformations

Brain modeling complements connectivity analysis by formalizing how epileptic dynamics emerge, propagate, and may be influenced under explicit dynamical assumptions. However, the practical question at this stage is not biological realism alone, but whether a model yields variables or summaries that can support localization, monitoring, planning, or constrained control under clinical and device constraints. Therefore, evaluation should focus on issues that determine real use: what is actually observed, which parameters can be identified from routine data, how much calibration is required, how outputs behave under drift, and what kinds of failure would mislead downstream decisions.[106]

We first consider models that can yield bounded state representations for monitoring and policy design, then planning-oriented models that can support target ranking or scenario exploration without serving as routine online control variables, and finally transformations between measurement spaces whose translational role remains narrower. Table 6 summarizes these model families along four deployment-oriented dimensions: dynamical commitment, observability, dominant failure modes, and translational role.

7.1 Deployment-relevant state representations for monitoring and policy design

The most clinically usable applications arise when models produce bounded state representations that can be estimated repeatedly, checked against data, and updated as conditions change.

Latent-state models. Latent-state models represent epileptic activity as transitions between hidden regimes and yield probabilistic state estimates that can support monitoring, triggering, and risk estimation. A seminal example cast seizure prediction within a stochastic three-state HMM comprising interictal, preictal, and ictal regimes and later implementations combined HMMs with time-frequency features for EEG-based detection and prediction.[107;108] Continuous state-space models provide a related alternative by tracking latent variables over time rather than restricting the system to discrete states. Heavy-tailed observation models can improve robustness to non-Gaussian artifacts and recent work has emphasized automated identification procedures aimed at reducing calibration burden and supporting scalable deployment.[109;110] Their practical value is straightforward: they compress noisy electrophysiology into a compact state estimate that can be thresholded, trended, recalibrated, or passed to downstream decision logic. However, that value depends on whether the inferred states remain interpretable and stable over time. State labels may drift, calibration may degrade, and operating thresholds may not transfer cleanly across recording conditions, patients, or longitudinal follow-up.

Reduced-order or surrogate control models. Low-dimensional surrogate models aim to preserve dominant epileptic dynamics while remaining simple enough for stability analysis and controller design. Such models have been used to illustrate how feedback may suppress pathological oscillations or steer trajectories away from seizure-prone regimes.[111;112] Data-driven reduced dynamics derived from intracranial recordings have similarly been proposed to support constrained stimulation design.[113] Their practical appeal is that they expose tractable control variables and explicit action trade-offs without requiring the full biological detail of a large mechanistic model. However, their validity is usually local to the regimes captured by the surrogate. Performance may deteriorate under drift, state change, sparse sampling, or changes in sensing and stimulation conditions. In practice, this limits their role mainly to controller design and policy exploration rather than autonomous long-term control.

7.2 Planning-oriented models rather than routine control objects

Greater biological detail does not by itself imply greater readiness for deployment. Therefore, more mechanistic and virtual frameworks are currently most useful when they help organize hypotheses, rank candidate interventions, or structure pre-intervention planning.

Mechanistic seizure models. Mechanistic seizure models seek to explain seizure onset, propagation, and termination as dynamical transitions in slow-fast systems, often linked to bifurcation structure.[114] The Epileptor family has become a canonical example, providing a flexible seizure-generating formalism across multiple dynamical regimes.[115;116] Neural-mass frameworks such as the Wendling family provide a complementary line of mechanistic modeling, linking epileptiform patterns more directly to mesoscopic excitatory-inhibitory circuit dynamics and recorded EEG or SEEG features.[117] Their appeal for intervention is that they introduce explicit latent quantities, such as excitability or slow permissivity-like processes, that suggest which aspects of the system an intervention would need to alter. Recent work has also begun to use simulation-based inference and related Bayesian procedures to fit mechanistic or virtual-brain models more directly to empirical data rather than relying on forward simulation alone.[118] However, these quantities are not directly observed and are rarely uniquely identifiable from routine clinical recordings. Better inference does not remove strong dependence on priors, model class, and recording completeness. Computational cost also remains non-trivial. As a result, these models can clarify mechanism and narrow the space of plausible interventions, but their practical role still depends on how reliably those latent quantities can be inferred from clinical data.

Network ictogenicity and virtual resection. Ictogenicity extends mechanistic reasoning from individual generators to networks by quantifying a configuration's propensity to generate seizures. Virtual resection studies use patient-specific networks to simulate removal of nodes or edges and thereby produce decision-oriented rankings of potentially influential regions.[119;120] This makes them attractive for candidate target ranking and pre-intervention planning. However, the rankings remain sensitive to connectome quality, incomplete sampling, and modeling assumptions. A high-ranking node in silico is not, by itself, sufficient evidence for intervention. Translational value therefore depends on whether the ranking adds information beyond existing clinical evidence and whether that added value holds in patient-specific validation.

Personalized virtual-brain frameworks. Personalized virtual-brain frameworks seek to combine patient-specific anatomy, connectivity, dynamics, and measurement mappings into a single in-silico environment.[121;122] Existing syntheses emphasize their potential for diagnosis, surgery planning, and stimulation strategy design.[123;124] The Virtual Epileptic Patient framework is among the most prominent examples, embedding epileptic dynamics within individual structural networks and fitting parameters to reproduce seizure spread.[125;126] More recent work has also shifted part of this literature from forward simulation toward parameter inference, including simulation-based and Bayesian approaches that aim to estimate patient-specific model parameters from empirical recordings.[118]

In principle, such frameworks can support pre-intervention exploration of candidate targets, stimulation timing, and parameter settings before clinical implementation.[111;113;115;127;128] Improved fitting is methodologically relevant because it addresses one of the main criticisms of virtual-brain modeling, namely the weak link between mechanistic variables and measured clinical data. However, the main deployment bottlenecks remain: incomplete observability, parameter non-uniqueness, prior sensitivity, computational burden, incomplete sampling, model mismatch, and unmodeled state shifts. These systems can still be useful in planning even when they are not fully identifiable, but that usefulness depends on explicit uncertainty reporting and patient-specific validation.

Taken together, these more mechanistic frameworks are currently strongest when used to constrain planning and interpretation, not when treated as routine online decision engines.

7.3 Transformations between measurement spaces

Transformations between measurement spaces matter because localization, planning, and intervention are often expressed in anatomical or source-level terms, whereas sensing is performed at the sensor level. Their value depends not on anatomical plausibility alone, but on whether they improve decision-relevant inference under realistic uncertainty.

Electrical source imaging. Electrical source imaging (ESI) can add prospective diagnostic value in presurgical evaluation, including in MRI-negative epilepsy.[129–131] Its usefulness lies in expressing event-related information at a more anatomically interpretable level than scalp topography alone. However, ESI remains highly sensitive to head-model assumptions, preprocessing choices, inverse regularization, and event selection. In deployment terms, the question is not whether the reconstruction looks anatomically plausible, but whether it improves localization or planning decisions beyond sensor-space analysis and does so reproducibly. Its value therefore lies in supporting localization and mapping under explicit uncertainty, not in directly recovering hidden ground truth.

Cross-modal and generative transformations. More recent approaches attempt to align or translate information across modalities, either by integrating multiple biomarkers into a unified source-space representation or by learning generative mappings between scalp and intracranial activity.[132;133] Related approaches include distributed spatial filtering for intracranial arrays.[134] These directions are conceptually attractive because they aim to narrow the gap between what is sensed and what is required for intervention planning. However, empirical constraints remain limited, and visually plausible reconstructions are not enough. These methods remain exploratory unless they can be shown to improve downstream localization, target ranking, or policy design under realistic acquisition limits. Concurrent scalp-intracranial recordings provide one of the few empirical routes for constraining such transformations by testing which sensor-level signatures correspond reproducibly to intracranial events in real recordings.[36]

8 Closed-Loop Neuromodulation and Online Adaptation

Closed-loop neuromodulation is the stage at which representation is translated into action. Sensed events, estimated risk states, network constraints, and model-derived summaries become relevant here only insofar as they can be mapped to stimulation decisions under explicit safety and device limits. Therefore, closed-loop systems are not simply detectors followed by stimulators, but constrained feedback architectures that must decide when, where, and how stimulation is delivered under partial observability, chronic non-stationarity, limited computation and power, and clinician oversight. From a translational perspective, the central question is whether the resulting decision logic remains safe, interpretable, and governable during long-term operation.

We first consider control variables and action design across reactive, risk-conditioned, and state-dependent paradigms. We then discuss adaptive algorithms and machine learning within governed control architectures, before turning to a control-theoretic framing across modalities and the practical deployment boundaries that define whether a policy can be regarded as clinically governable (Table 7).

8.1 Control variables and action design across closed-loop paradigms

Closed-loop paradigms differ less by algorithmic label than by which variable is used for control and which part of the action space is allowed to adapt. Reactive systems use detected abnormal activity as the control variable and typically adapt timing within a tightly bounded stimulation programme. Risk-conditioned approaches instead use slower estimates of seizure likelihood or contextual covariates to influence scheduling or operating mode. More elaborate frameworks may incorporate latent states or network-derived constraints to guide contact selection, parameter choice, or mode switching. Framed this way, reactive, predictive, and adaptive systems are not separate literatures so much as different policy designs acting on different state summaries. Therefore, the practical issue is not only how neural activity is represented, but how those representations are operationalized into auditable and reproducible decision rules within clinician-facing workflows.[4;135]

Existing clinical neuromodulation already shows that practical success depends on patient selection, conservative programming, longitudinal adjustment, and tolerability as well as on stimulation delivery itself.[136] Clinical experience with epilepsy DBS, including centromedian thalamic targets, and pediatric intracranial stimulation similarly indicates that long term outcomes are syndrome-dependent and remain strongly shaped by programming choices and tolerability constraints.[137;138] Responsive neurostimulation (RNS) operationalized the canonical sensing-detection-stimulation loop for focal epilepsy.[139] As a control architecture, it exemplifies a reactive policy in which detected abnormal activity triggers tightly bounded stimulation bursts. This reactive paradigm sits alongside other approved neuromodulation therapies, including vagal nerve stimulation (VNS) and DBS, whose distinct targets and stimulation schedules define the current clinical baseline for closed-loop development.[136] Long-term follow-up shows that benefit can be durable over years, but optimization depends on repeated clinical reprogramming of detection features, electrode selection, and stimulation parameters.[140;141] Therefore, practical RNS optimization is dominated by clinician-led personalization over longitudinal follow-up rather than by autonomous online learning.[142] Heterogeneity across patients and substrates, including mesial temporal lobe epilepsy, further shows that deployment success depends on calibration, conservative operating points, and programming strategy as much as on the original detector design.[143]

DBS of the anterior nucleus of the thalamus (ANT) provides a complementary reference. Although typically delivered more continuously, DBS-ANT demonstrates that seizure reduction can arise through modulation of distributed networks rather than focal event suppression alone.[144] This widens the design space for closed-loop policies: action may involve not only immediate triggering, but also scheduling, parameter selection, or clinician-guided mode changes. Comparative reviews across VNS, DBS, and RNS reinforce this point by showing that differences in targets and stimulation paradigms primarily alter opportunities for personalization and adaptation rather than defining wholly distinct control principles.[145]

The key design question is whether the control variable is stable enough to govern an action safely over time. Event-triggered policies remain the most mature because their action space is narrow and their operational logic is explicit. As policies move toward risk conditioning, network-informed constraints, or state-dependent mode switching, the burden shifts toward calibration, monitoring, and governance rather than toward estimator sophistication alone. Recent exploratory work on stimulating ictal fast-ripple hubs makes this trade-off concrete: network-informed targets may improve efficacy and reduce seizure-evocation risk relative to direct seizure onset zone (SOZ) stimulation, but they also depend on high-frequency biomarkers whose online detection and longitudinal stability remain difficult to guarantee.[146] Therefore, such approaches are conceptually relevant, but not yet validated as robust long-term control variables for routine deployment.

8.2 Adaptive algorithms, machine learning, and governability

Adaptive neuromodulation generalizes reactive triggering by allowing stimulation parameters, operating modes, or control thresholds to vary with estimated neural state. Closed-loop adaptation also operates over multiple time scales, including rapid within-session reactions over milliseconds to seconds, clinician-driven parameter tuning over days to weeks, and slower policy updates that redefine states, thresholds, or decision rules over weeks to months. Current implanted epilepsy systems mainly operate at the first two time scales.

Against this background, adaptive DBS in Parkinson's disease provides a useful template; a biomarker defines a low-dimensional state estimate, decision logic maps that state to bounded actions, and stimulation is delivered under explicit constraints.[147] Implantable deployments in that field also make clear why algorithmic complexity is not the primary design axis. Power consumption, latency, signal quality, calibration burden, and long-term operability all favor simple and interpretable state representations together with bounded policies.[148;149]

In epilepsy, machine learning is presently most defensible when used for state estimation rather than unconstrained policy learning. Robust biomarker estimation, latent-state tracking, and bounded mode switching within clinician-defined limits are all plausible roles. By contrast, candidate biomarkers for adaptive DBS in epilepsy have not yet matured into sufficiently validated control variables with prospective evidence that closed-loop adjustment improves outcomes.[150;151] There is a meaningful difference between using learning to support a governed control surface and using it to revise policy autonomously without reliable safeguards.

This is why reinforcement learning and more autonomous adaptive schemes should currently be interpreted cautiously. The main barrier is not reward design in isolation, but the combination of partial observability, sparse and delayed outcomes, hard safety envelopes, a controlled system that itself changes over time, and the requirement for clinician supervision. In this setting, the relevant question is whether adaptation can be bounded, audited, and clinically governed. Therefore, greater policy flexibility should not be assumed to imply greater translational value. The most credible near-term direction is governed adaptation built around stable state summaries, bounded action spaces, and explicit clinician supervision.

8.3 Control-theoretic framing across modalities

Across stimulation modalities, closed-loop neuromodulation is most coherently understood as constrained feedback control of a dynamical system in which stimulation is the control input and neural recordings are noisy, partial observations of the controlled state. This framing helps distinguish differences in implementation and clinical objective from differences in underlying control principle.

In practice, closed-loop design is dominated by constraint handling. Stability, robustness, safety limits, and tolerance to drift are primary objectives and often outweigh theoretical optimality.[152] Conceptual analyses likewise emphasize that stimulation effects are network and state-dependent, cautioning against simplistic mechanistic narratives and motivating context-aware policies.[153] Experimental evidence supports this view: weak but physiologically plausible electric fields can entrain ongoing neocortical activity in a manner that depends on the current dynamical regime.[154] Therefore, the implication for control design is practical rather than abstract: policies should be built around explicit state-conditional biomarkers, bounded action envelopes, and recalibration paths that remain workable under drift.

8.4 Deployment boundaries for closed-loop policies

A closed-loop policy is not deployment-ready unless its action space is bounded, its failure modes are explicit, conservative fallback behavior exists under uncertainty or signal loss, updates are logged and reversible, and clinician authority is preserved. Therefore, end-to-end latency, computational and power budget, charge density and duty-cycle limits, drift monitoring, and the clinician-facing control surface should be treated as part of the policy itself rather than as reporting formalities. In practice, these conditions define the minimum boundary between a governable clinical system and a research prototype.

9 Open Challenges and Future Directions

Despite progress across sensing, modeling, and control, major unresolved challenges remain at the level of long-term deployment. The remaining obstacles are not only algorithmic, but also methodological and clinical: how trustworthy operating points should be defined and maintained, how adaptive policies should be governed over time, and what level of evidence is sufficient before technical outputs are allowed to influence care.

9.1 Evaluation, longitudinal robustness, and uncertainty

A central open challenge is how to define and maintain a trustworthy operating point within patients over time. The unresolved issue is not simply whether performance degrades under chronic drift, but how such degradation should be detected, measured, and acted upon once patient-specific variability, state dependence, and chronobiological structure are all present.[155] Current evaluation practice still relies heavily on short retrospective settings, so strong offline results often remain a weak proxy for false alarm burden, calibration stability, and operational reliability during prolonged use.[156]

This challenge has several practical components. It remains unclear how long-term systems should detect loss of robustness, decide when recalibration is required, and distinguish tolerable variation from clinically important degradation. It is similarly unclear which validation designs should be treated as minimally credible for longitudinal deployment, particularly when thresholds, prevalence, and background dynamics evolve within patients rather than only across cohorts.

Under these conditions, uncertainty is not a secondary refinement, but part of the deployment interface itself. If a system cannot represent when its outputs have become unreliable, then apparent continuity of operation may conceal clinically important failure. Therefore, uncertainty estimation, abstention mechanisms, and patient-level monitoring of calibration drift should be incorporated into deployment-facing validation rather than treated as optional additions. Longer prospective recordings, explicit control of temporal leakage, and benchmark designs that test preservation of a trustworthy operating point during prolonged use remain key methodological needs. [156]

9.2 Control safety and adaptive governance

A second open challenge is how adaptive closed-loop systems can remain governable as conditions change over time. Safety limits how far autonomy can be extended, yet static policies may fail as signals, states, and treatment effects evolve. Therefore, the key question is not simply whether adaptation is desirable, but which forms of adaptation can be permitted without undermining validation, oversight, and rollback.

What remains unresolved is how safe adaptation should be demonstrated once policies, thresholds, or state definitions are allowed to change during prolonged use, particularly under partial observability and changing treatment response. The practical challenge is to define governance structures that specify which updates may occur automatically, which require clinician approval, how escalation and rollback should be triggered, and how failure under drift, signal loss, or conflicting evidence should be handled. Therefore, the immediate methodological need is for prospective protocols for supervised adaptation, explicit criteria for escalation and rollback, and validation designs that test controlled degradation under clinically realistic failure modes as well as nominal benefit.

9.3 Clinical integration, evidential sufficiency, and deployment value

A further open challenge is how clinical value should be demonstrated before technical outputs are allowed to influence care. Algorithm outputs must reduce burden or improve decisions in practice, not merely improve retrospective metrics. However, workflow benefit remains defined and measured inconsistently. Effects on review time, prioritization, agreement, programming decisions, tolerance of uncertainty, and handling of failure are still much less standardized than conventional classification endpoints.

The unresolved issue is not only whether a signal, biomarker, or model output is informative, but what level of evidence is sufficient before it can influence clinical decision-making. In practice, the field still lacks widely adopted evidential tiers that distinguish exploratory markers from decision-support tools, and decision-support tools from outputs robust enough for bounded intervention use. These tiers should make explicit not only expected benefit, but also uncertainty, operational consequence, and acceptable failure.

Therefore, progress will be most convincing when technical gains are tied to clinician-centered endpoints and prospective evidence of decision impact. Clearer evidential standards should distinguish exploratory outputs from those robust enough to support workflow, programming, or bounded intervention use.

10 Conclusion

This review examines epilepsy-related methods relevant to algorithm-guided neuromodulation through a deployment-oriented lens, treating each pipeline stage from data acquisition through preprocessing, event and evoked-response sensing, connectivity analysis, brain modeling, and closed-loop control as part of a tightly coupled system in which upstream choices constrain downstream reliability. Three arguments structure this synthesis.

First, neuromodulation is best understood as a constrained control problem under partial observability and non-stationarity, not simply as a sequence of detection tasks. Both diagnostic probing (via SPES-evoked responses) and therapeutic intervention (via RNS, DBS, and related modalities) act within patient-specific networks that themselves change over time through stimulation-induced plasticity. Algorithms that ignore these dynamics, optimizing for offline accuracy without accounting for drift, calibration decay, or state dependence, often fail to remain reliable when deployed chronically.

Second, several foundational debates remain unresolved, and their resolution will shape the next generation of neuromodulation systems. Whether reliable seizure prediction is achievable, or whether we are limited to probabilistic risk modulation based on circadian and multidien cycles, determines how far anticipatory intervention can extend. Whether targeting should emphasize focal seizure onset or network-level properties such as hub structures and propagation pathways determines how connectivity informs programming. Whether candidate biomarkers derived from event sensing, connectivity analysis, and latent-state estimation are sufficiently validated to serve as closed-loop control variables determines how far adaptation can safely proceed. This review has examined the available evidence on each debate, but definitive answers will require prospective validation under deployment-realistic conditions that most current studies do not yet provide.

Third, progress toward clinically durable systems depends less on marginal algorithmic gains than on evaluation reform and explicit robustness requirements. We have emphasized stageappropriate reporting and evaluation practices, prioritizing metrics that reflect clinical burden and longitudinal reliability, including false positives per unit time, time-in-warning, calibration stability, and early-window fidelity, rather than retrospective accuracy on curated datasets. Cross-method comparison also remains limited by heterogeneous datasets, inconsistent evaluation protocols, and incomplete reporting, a limitation that benchmark development and standardization efforts must address.

Looking forward, the most credible path to clinically acceptable closed-loop neuromodulation lies not in fully autonomous optimization, but in governed adaptation: bounded policies built around stable state summaries, explicit safety constraints, clinician-supervised recalibration, and auditable update mechanisms. The framework and shared coordinate system proposed here, organizing methods by temporal scale, observability, drift vulnerability, and deployment role, provide a shared vocabulary for interpreting heterogeneous methods in deployment terms and for identifying where current evidence is strongest and where critical gaps remain.

Ultimately, the goal is not algorithmic sophistication for its own sake, but systems that remain safe, interpretable, and effective over the years-long time scales that matter to patients and clinicians. Achieving that goal will require not only technical advances, but also a sustained shift in how the field evaluates, reports, and validates its methods.

List of Abbreviations

AI	Artificial intelligence
ANT	Anterior nucleus of the thalamus
ASR	Artifact subspace reconstruction
AUC	Area under the curve
BIDS	Brain Imaging Data Structure
BPC	Basis profile curve
BSS	Blind source separation
CCEP	Cortico-cortical evoked potential
CHB-MIT	Children's Hospital Boston-MIT (EEG database)
cEI	Connectivity epileptogenicity index
CRP	Canonical response parameterization
DBS	Deep brain stimulation
DBS-ANT	Deep brain stimulation of the anterior nucleus of the thalamus
DR	Delayed response
DRTE	Digital Response Test in Epilepsy
DTI	Diffusion tensor imaging
EC	Effective connectivity
EEG	Electroencephalography
ECoG	Electrocorticography
ER	Early response
ESI	Electrical source imaging
FC	Functional connectivity
fMRI	Functional magnetic resonance imaging
FP/min	False positives per minute
HMM	Hidden Markov model
ICA	Independent component analysis
IED	Interictal epileptiform discharge

iEEG	Intracranial EEG
iEEG-BIDS	Intracranial EEG extension of the Brain Imaging Data Structure
IEEE	Institute of Electrical and Electronics Engineers
MEG	Magnetoencephalography
ML	Machine learning
MPARRM	Multi-channel/pulse artifact removal by reconstruction method
MRI	Magnetic resonance imaging
N1	Negative component 1
PET	Positron emission tomography
PPV	Positive predictive value
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
ROC	Receiver operating characteristic
RNS	Responsive neurostimulation
SC	Structural connectivity
scEEG	Scalp EEG
sEEG	Stereo-electroencephalography
Sens/Prec	Sensitivity/Precision
SNR	Signal-to-noise ratio
SOZ	Seizure onset zone
SPES	Single-pulse electrical stimulation
SSM	State-space model
SSP	Signal-space projection
SVM	Support vector machine
TUH	Temple University Hospital
VNS	Vagus nerve stimulation
wPLI	Weighted phase lag index

Conflict of Interest

The authors declare no conflicts of interest.

Financial Disclosure

This work received no specific funding.

Data Availability Statement

No new data were generated or analysed in support of this review. All information discussed is derived from previously published studies, which are cited in the manuscript.

Data Sharing Statement

This article is a review and does not report primary datasets. Therefore, no new data are available for sharing. Any datasets referenced are available from their original sources as cited.

References

1. Moutonnet N, White S, Campbell BP, et al. Clinical translation of machine learning algorithms for seizure detection in scalp electroencephalography: systematic review. arXiv preprint arXiv:2404.15332. 2024.
2. Lucas A, Revell A, Davis KA. Artificial intelligence in epilepsy-applications and pathways to the clinic. Nat Rev Neurol. 2024;20(6):319–336.

3. Kerr WT, McFarlane KN, Figueiredo Pucci G. The present and future of seizure detection, prediction, and forecasting with machine learning, including the future impact on clinical trials. Front Neurol. 2024;15:1425490.

4. Daraie AH, Damiani A, Khoshkhou M, et al. Artificial intelligence for adaptive neuromodulation in drug-resistant epilepsy. Epilepsia. 2026:1–16.

5. Jobst BC, Bartolomei F, Diehl B, et al. Intracranial EEG in the 21st century. Epilepsy Curr. 2020;20(4):180–188.

6. Ung H, Baldassano SN, Bink H, et al. Intracranial EEG fluctuates over months after implanting electrodes in human brain. J Neural Eng. 2017;14(5):056011.

7. Viana PF, Duun-Henriksen J, Biondi A, et al. Real-world epilepsy monitoring with ultra-long-term subcutaneous electroencephalography: a 15-month prospective study. Epilepsia. 2025;66(11):4476–4489.

8. Biondi A, Simblett SK, Viana PF, et al. Feasibility and acceptability of an ultra-long-term at-home EEG monitoring system (EEG@ HOME) for people with epilepsy. Epilepsy Behav. 2024;151:109609.

9. Macea J, Bhagubai M, Broux V, De Vos M, Van Paesschen W. In-hospital and home-based long-term monitoring of focal epilepsy with a wearable electroencephalographic device: diagnostic yield and user experience. Epilepsia. 2023;64(4):937–950.

10. Obeid I, Picone J. The Temple University Hospital EEG data corpus. Front Neurosci. 2016;10:196.

11. Shah V, Von Weltin E, Lopez S, et al. The Temple University Hospital seizure detection corpus. Front Neuroinform. 2018;12:83.

12. Goldberger AL, Amaral LAN, Glass L, et al. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. Circulation. 2000;101(23):E215–E220.

13. Sun C, Jing J, Turley N, et al. Harvard Electroencephalography Database: a comprehensive clinical electroencephalographic resource from four Boston hospitals. Epilepsia. 2025;66(9):3411–3425.
14. Klatt J, Feldwisch-Drentrup H, Ihle M, et al. The EPILEPSIAE database: an extensive electroencephalography database of epilepsy patients. 2012.
15. Wong S, Simmons A, Rivera-Villicana J, et al. EEG datasets for seizure detection and prediction-a review. Epilepsia Open. 2023;8(2):252–267.
16. Cook MJ, O'Brien TJ, Berkovic SF, et al. Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: a first-in-man study. Lancet Neurol. 2013;12(6):563–571.
17. Karoly PJ, Ung H, Grayden DB, et al. The circadian profile of epilepsy improves seizure forecasting. Brain. 2017;140(8):2169–2182.
18. Lin N, Zheng M, Li L, et al. An EEG dataset for interictal epileptiform discharge with spatial distribution information. Sci Data. 2025;12(1):229.
19. van Blooijs D, Blok S, Huiskamp GJM, Leijten FSS. PRIOS. 2023.
20. Kini LG, Davis KA, Wagenaar JB. Data integration: combined imaging and electrophysiology data in the cloud. Neuroimage. 2016;124:1175–1181.
21. Holdgraf C, Appelhoff S, Bickel S, et al. iEEG-BIDS, extending the Brain Imaging Data Structure specification to human intracranial electrophysiology. Sci Data. 2019;6(1):102.
22. Sharma R, Meena HK. Emerging trends in EEG signal processing: a systematic review. SN Comput Sci. 2024;5(4):415.
23. Chaddad A, Wu Y, Kateb R, Bouridane A. Electroencephalography signal processing: a comprehensive review and analysis of methods and techniques. Sensors (Basel). 2023;23(14):6434.
24. Zhang X, Zhang X, Huang Q, Chen F. A review of epilepsy detection and prediction methods based on EEG signal processing and deep learning. Front Neurosci. 2024;18:1468967.

25. Djemal A, Bouchaala D, Fakhfakh A, Kanoun O. Artifacts removal from epileptic EEG signal based on independent components analysis method. In: 2022 IEEE International Symposium on Medical Measurements and Applications (MeMeA). 2022:1–6.
26. Lopes F, Leal A, Medeiros J, et al. Automatic electroencephalogram artifact removal using deep convolutional neural networks. IEEE Access. 2021;9:149955–149970.
27. Wu YH, Lin HC, Huang CW, Wu CY, Ker MD. Stimulation-induced artifact removal of the local field potential through hardware design: toward the implantable closed-loop deep brain stimulation. IEEE Access. 2024.
28. Zhou A, Johnson BC, Muller R. Toward true closed-loop neuromodulation: artifact-free recording during stimulation. Curr Opin Neurobiol. 2018;50:119–127.
29. Varone G, Hussain Z, Sheikh Z, et al. Real-time artifacts reduction during TMS-EEG co-registration: a comprehensive review on technologies and procedures. Sensors (Basel). 2021;21(2):637.
30. Trebaul L, Rudrauf D, Job AS, et al. Stimulation artifact correction method for estimation of early cortico-cortical evoked potentials. J Neurosci Methods. 2016;264:94–102.
31. Xie T, Foutz TJ, Adamek M, et al. Single-pulse electrical stimulation artifact removal using the novel matching pursuit-based artifact reconstruction and removal method (MPARRM). J Neural Eng. 2023;20(6):066036.
32. Lim J, Wang PT, Bashford L, et al. Suppression of cortical electrostimulation artifacts using pre-whitening and null projection. J Neural Eng. 2023;20(5):056018.
33. Yan X, Boudrias MH, Mitsis GD. Removal of transcranial alternating current stimulation EEG artifacts using blind source separation and wavelets. IEEE Trans Biomed Eng. 2022;69(10):3183–3192.
34. Fernandez-de-Retana M, Matanzas-de-Luis P, Peña J, Almeida A. A deep learning approach to artifact removal in transcranial electrical stimulation: from shallow methods to deep neural networks and state space models. Neuroscience. 2025.

35. Halford JJ. Computerized epileptiform transient detection in the scalp electroencephalogram: obstacles to progress and the example of computerized ECG interpretation. Clin Neurophysiol. 2009;120(11):1909–1915.
36. Spyrou L, Martín-Lopez D, Valentín A, Alarcón G, Sanei S. Detection of intracranial signatures of interictal epileptiform discharges from concurrent scalp EEG. Int J Neural Syst. 2016;26(04):1650016.
37. Sudhakar JG, Elango K. A comprehensive review of EEG-based seizure detection techniques. IEEE Access. 2025.
38. Kuhlmann L, Lehnertz K, Richardson MP, Schelter B, Zaveri HP. Seizure prediction-ready for a new era. Nat Rev Neurol. 2018;14(10):618–630.
39. Rasheed K, Qayyum A, Qadir J, et al. Machine learning for predicting epileptic seizures using EEG signals: a review. IEEE Rev Biomed Eng. 2020;14:139–155.
40. Andrzejak RG, Zaveri HP, Schulze-Bonhage A, et al. Seizure forecasting: where do we stand? Epilepsia. 2023;64:S62–S71.
41. Wu Y, Su ELM, Wu M, Ooi CY, Holderbaum W. A review of machine learning and deep learning trends in EEG-based epileptic seizure prediction. IEEE Access. 2025.
42. Carmo AS, Abreu M, Baptista MF, et al. Automated algorithms for seizure forecast: a systematic review and meta-analysis. J Neurol. 2024;271(10):6573–6587.
43. Pal Attia T, Viana PF, Nasseri M, et al. Seizure forecasting using minimally invasive, ultra-long-term subcutaneous EEG: generalizable cross-patient models. Epilepsia. 2023;64:S114–S123.
44. Viana PF, Pal Attia T, Nasseri M, et al. Seizure forecasting using minimally invasive, ultra-long-term subcutaneous electroencephalography: individualized intra-patient models. Epilepsia. 2023;64:S124–S133.
45. Yang H, Mueller J, Eberlein M, et al. Seizure forecasting with ultra long-term EEG signals. Clin Neurophysiol. 2024;167:211–220.

46. Sun Y, Guan M, Chen X, et al. Deep learning-based classification and segmentation of interictal epileptiform discharges using multichannel electroencephalography. Epilepsia. 2025.

47. Tong PF, Dong B, Zeng X, Chen L, Chen SX. Detection of interictal epileptiform discharges using transformer based deep neural network for patients with self-limited epilepsy with centrotemporal spikes. Biomed Signal Process Control. 2025;101:107238.

48. Lin N, Li L, Gao W, et al. Development and validation of a multimodal automatic interictal epileptiform discharge detection model: a prospective multi-center study. BMC Med. 2025;23(1):479.

49. Wu Y, Lu L, Xu A, et al. Neural networks for epilepsy detection and prediction with EEG signals: a systematic review. Artif Intell Rev. 2026;59(1):30.

50. Ein Shoka AA, Dessouky MM, El-Sayed A, Hemdan EED. EEG seizure detection: concepts, techniques, challenges, and future trends. Multimed Tools Appl. 2023;82(27):42021–42051.

51. Xiang F, Liu M, Chen W, Zheng S, Zhang J, Du G. A deep hybrid CSAE-GRU framework with two-stage balancing for automatic epileptic seizure detection using EEG-derived features. Front Neurosci. 2025;19:1698960.

52. Zhu R, Pan WX, Liu JX, Shang JL. Epileptic seizure prediction via multidimensional transformer and recurrent neural network fusion. J Transl Med. 2024;22(1):895.

53. Koutsouvelis P, Chybowski B, Gonzalez-Sulser A, Abdullateef S, Escudero J. Preictal period optimization for deep learning-based epileptic seizure prediction. J Neural Eng. 2024;21(6):066040.

54. Karoly PJ, Cook MJ, Maturana M, et al. Forecasting cycles of seizure likelihood. Epilepsia. 2020;61(4):776–786.

55. Yang H, Müller J, Kalousios S, et al. Seizure forecasting with epilepsy cycles: on the causality of forecasting pipelines. Epilepsia. 2026.

56. Proix T, Truccolo W, Leguia MG, et al. Forecasting seizure risk in adults with focal epilepsy: a development and validation study. Lancet Neurol. 2021;20(2):127–135.

57. Stirling RE, Cook MJ, Grayden DB, Karoly PJ. Seizure forecasting and cyclic control of seizures. Epilepsia. 2021;62:S2–S14.

58. Shafiezadeh S, Duma GM, Pozza M, Testolin A. A systematic review of cross-patient approaches for EEG epileptic seizure prediction. J Neural Eng. 2024;21(6):061004.

59. von Allmen A, Lu D, Jagella C, et al. Digital Response Test in Epilepsy assesses interictal epileptiform discharge effects in real time. Epilepsia. 2025.

60. Pinto MF, Batista J, Leal A, et al. The goal of explaining black boxes in EEG seizure prediction is not to explain models' decisions. Epilepsia Open. 2023;8(2):285–297.

61. Saboo KV, Cao Y, Kremen V, et al. Individualized seizure cluster prediction using machine learning and chronic ambulatory intracranial EEG. IEEE Trans Nanobioscience. 2023;22(4):818–827.

62. Matsumoto R, Nair DR, LaPresto E, et al. Functional connectivity in the human language system: a cortico-cortical evoked potential study. Brain. 2004;127(10):2316–2330.

63. Keller CJ, Honey CJ, Mégevand P, Entz L, Ulbert I, Mehta AD. Mapping human brain networks with cortico-cortical evoked potentials. Philos Trans R Soc Lond B Biol Sci. 2014;369(1653):20130528.

64. Prime D, Rowlands D, O'Keefe S, Dionisio S. Considerations in performing and analyzing the responses of cortico-cortical evoked potentials in stereo-EEG. Epilepsia. 2018;59(1):16–26.

65. Hajnal B, Szabó JP, Tóth E, et al. Intracortical mechanisms of single pulse electrical stimulation (SPES) evoked excitations and inhibitions in humans. Sci Rep. 2024;14(1):13784.

66. Qiang Z, Norris J, Cooray G, et al. Cortico-cortical evoked potentials: analytical techniques and emerging paradigms for epileptogenic zone localization. Epilepsia. 2025.

67. Valentín A, Alarcón G, Honavar M, et al. Single pulse electrical stimulation for identification of structural abnormalities and prediction of seizure outcome after epilepsy surgery: a prospective study. Lancet Neurol. 2005;4(11):718–726.
68. Flanagan D, Valentín A, García Seoane JJ, Alarcón G, Boyd SG. Single-pulse electrical stimulation helps to identify epileptogenic cortex in children. Epilepsia. 2009;50(7):1793–1803.
69. Nayak D, Valentín A, Selway RP, Alarcón G. Can single pulse electrical stimulation provoke responses similar to spontaneous interictal epileptiform discharges? Clin Neurophysiol. 2014;125(7):1306–1311.
70. Iwasaki M, Enatsu R, Matsumoto R, et al. Accentuated cortico-cortical evoked potentials in neocortical epilepsy in areas of ictal onset. Epileptic Disord. 2010;12(4):292–302.
71. Enatsu R, Piao Z, O'Connor T, et al. Cortical excitability varies upon ictal onset patterns in neocortical epilepsy: a cortico-cortical evoked potential study. Clin Neurophysiol. 2012;123(2):252–260.
72. Trebaul L, Deman P, Tuyisenge V, et al. Probabilistic functional tractography of the human cortex revisited. Neuroimage. 2018;181:414–429.
73. Cornblath EJ, Lucas A, Armstrong C, et al. Quantifying trial-by-trial variability during cortico-cortical evoked potential mapping of epileptogenic tissue. Epilepsia. 2023;64(4):1021–1034.
74. Al-Sadek T, Wadhwa A, Wadhwa M, Warren AEL, Rolston JD. Methodologies to detect cortico-cortical evoked potentials: a systematic review. Front Hum Neurosci. 2025;19:1636115.
75. Miller KJ, Müller KR, Hermes D. Basis profile curve identification to understand electrical stimulation effects in human brain networks. PLoS Comput Biol. 2021;17(9):e1008710.
76. Miller KJ, Müller KR, Valencia GO, et al. Canonical response parameterization: quantifying the structure of responses to single-pulse intracranial electrical brain stimulation. PLoS Comput Biol. 2023;19(5):e1011105.

77. Kamali G, Smith RJ, Hays M, et al. Transfer function models for the localization of seizure onset zone from cortico-cortical evoked potentials. Front Neurol. 2020;11:579961.

78. Zhao C, Liang Y, Li C, et al. Localization of epileptogenic zone based on cortico-cortical evoked potential (CCEP): a feature extraction and graph theory approach. Front Neuroinform. 2019;13:31.

79. Hays MA, Smith RJ, Wang Y, et al. Cortico-cortical evoked potentials in response to varying stimulation intensity improves seizure localization. Clin Neurophysiol. 2023;145:119–128.

80. Yang B, Zhao B, Li C, et al. Localizing seizure onset zone by a cortico-cortical evoked potentials-based machine learning approach in focal epilepsy. Clin Neurophysiol. 2024;158:103–113.

81. van den Boom MA, Gregg NM, Valencia GO, et al. ER-detect: a pipeline for robust detection of early evoked responses in BIDS-iEEG electrical stimulation data. J Neurosci Methods. 2025;418:110389.

82. Patel SA, Brinyark H, Coyne C, et al. Cortico-cortical evoked potentials: automated localization and classification of early and late responses. J Neurosci Methods. 2025:110571.

83. Dou Y, Xia J, Fu M, Cai Y, Meng X, Zhan Y. Identification of epileptic networks with graph convolutional network incorporating oscillatory activities and evoked synaptic responses. Neuroimage. 2023;284:120439.

84. Friston KJ. Functional and effective connectivity: a review. Brain Connect. 2011;1(1):13–36.

85. Yaffe RB, Borger P, Megevand P, et al. Physiology of functional and effective networks in epilepsy. Clin Neurophysiol. 2015;126(2):227–236.

86. Parker CS, Clayden JD, Cardoso MJ, et al. Structural and effective connectivity in focal epilepsy. Neuroimage Clin. 2018;17:943–952.

87. Nolte G, Bai O, Wheaton L, Mari Z, Vorbach S, Hallett M. Identifying true brain interaction from EEG data using the imaginary part of coherency. Clin Neurophysiol. 2004;115(10):2292–2307.

88. Vinck M, Oostenveld R, Van Wingerden M, Battaglia F, Pennartz CMA. An improved index of phase-synchronization for electrophysiological data in the presence of volume conduction, noise and sample-size bias. Neuroimage. 2011;55(4):1548–1565.

89. Nagy P, Tóth B, Winkler I, Boncz Á. The effects of spatial leakage correction on the reliability of EEG-based functional connectivity networks. Hum Brain Mapp. 2024;45(8):e26747.

90. Van Mierlo P, Höller Y, Focke NK, Vulliemoz S. Network perspectives on epilepsy using EEG/MEG source connectivity. Front Neurol. 2019;10:721.

91. Bernhardt BC, Bonilha L, Gross DW. Network analysis for a network disorder: the emerging role of graph theory in the study of epilepsy. Epilepsy Behav. 2015;50:162–170.

92. Royer J, Bernhardt BC, Larivière S, et al. Epilepsy and brain network hubs. Epilepsia. 2022;63(3):537–550.

93. Rijal S, Corona L, Perry MS, et al. Functional connectivity discriminates epileptogenic states and predicts surgical outcome in children with drug resistant epilepsy. Sci Rep. 2023;13(1):9622.

94. Balatskaya A, Roehri N, Lagarde S, et al. The “Connectivity Epileptogenicity Index” (cEI), a method for mapping the different seizure onset patterns in Stereo ElectroEncephalography recorded seizures. Clin Neurophysiol. 2020;131(8):1947–1955.

95. Liu F, Wang Y, Li M, et al. Dynamic functional network connectivity in idiopathic generalized epilepsy with generalized tonic-clonic seizure. Hum Brain Mapp. 2017;38(2):957–973.

96. Cousyn L, Messaoud RB, Lehongre K, et al. Daily resting-state intracranial EEG connectivity for seizure risk forecasts. Epilepsia. 2023;64(2):e23–e29.

97. Qin L, Zhou Q, Sun Y, Pang X, Chen Z, Zheng J. Dynamic functional connectivity and gene expression correlates in temporal lobe epilepsy: insights from hidden Markov models. J Transl Med. 2024;22(1):763.

98. Piper RJ, Richardson RM, Worrell G, et al. Towards network-guided neuromodulation for epilepsy. Brain. 2022;145(10):3347–3362.
99. Mehta K, Damiani A, Pirondini E, Agashe S, McIntyre CC, Gonzalez-Martinez JA. Leveraging functional and structural connectomics to guide neuromodulation in epilepsy. J Clin Neurophysiol. 2025;42(6):521–526.
100. Charlebois CM, Anderson DN, Johnson KA, et al. Patient-specific structural connectivity informs outcomes of responsive neurostimulation for temporal lobe epilepsy. Epilepsia. 2022;63(8):2037–2055.
101. Ji GJ, Fox MD, Morton-Dutton M, et al. A generalized epilepsy network derived from brain abnormalities and deep brain stimulation. Nat Commun. 2025;16(1):2783.
102. Kobayashi K, Taylor KN, Shahabi H, et al. Effective connectivity relates seizure outcome to electrode placement in responsive neurostimulation. Brain Commun. 2024;6(1):fcae035.
103. Hu X, Yao Y, Zhao B, et al. Effective connectivity predicts surgical outcomes in temporal lobe epilepsy: a SEEG study. CNS Neurosci Ther. 2025;31(8):e70563.
104. Gregg NM, Ojeda Valencia G, Pridalova T, et al. Thalamic stimulation induced changes in network connectivity and excitability in epilepsy. Ann Neurol. 2025.
105. Matarrese MA, Loppini A, Fabbri L, et al. Spike propagation mapping reveals effective connectivity and predicts surgical outcome in epilepsy. Brain. 2023;146(9):3898–3912.
106. Stefanescu RA, Shivakeshavan RG, Talathi SS. Computational models of epilepsy. Seizure. 2012;21(10):748–759.
107. Wong S, Gardner AB, Krieger AM, Litt B. A stochastic framework for evaluating seizure prediction algorithms using hidden Markov models. J Neurophysiol. 2007;97(3):2525–2532.
108. Abdullah MH, Abdullah JM, Abdullah MZ. Seizure detection by means of hidden Markov model and stationary wavelet transform of electroencephalograph signals. In: Proceedings of the 2012 IEEE-EMBS International Conference on Biomedical and Health Informatics. 2012:62–65.

109. Wang Y, Qi Y, Zhu J, et al. A cauchy-based state-space model for seizure detection in EEG monitoring systems. IEEE Intell Syst. 2014;30(1):6–12.

110. Wang Z, Sperling MR, Wyeth D, Guez A. Automated seizure detection based on state-space model identification. Sensors (Basel). 2024;24(6):1902.

111. Ge Y, Cao Y, Yi G, et al. Robust closed-loop control of spike-and-wave discharges in a thalamocortical computational model of absence epilepsy. Sci Rep. 2019;9(1):9093.

112. Chakravarthy N, Sabesan S, Tsakalis K, Iasemidis L. Controlling epileptic seizures in a neural mass model. J Comb Optim. 2009;17(1):98–116.

113. Ashourvan A, Pequito S, Khambhati AN, et al. Model-based design for seizure control by stimulation. J Neural Eng. 2020;17(2):026009.

114. Jirsa VK, Stacey WC, Quilichini PP, Ivanov AI, Bernard C. On the nature of seizure dynamics. Brain. 2014;137(8):2210–2230.

115. El Houssaini K, Ivanov AI, Bernard C, Jirsa VK. Seizures, refractory status epilepticus, and depolarization block as endogenous brain activities. Phys Rev E Stat Nonlin Soft Matter Phys. 2015;91(1):010701.

116. El Houssaini K, Bernard C, Jirsa VK. The epileptor model: a systematic mathematical analysis linked to the dynamics of seizures, refractory status epilepticus, and depolarization block. eNeuro. 2020;7(2).

117. Wendling F, Koksal-Ersoz E, Al-Harrach M, et al. Multiscale neuro-inspired models for interpretation of EEG signals in patients with epilepsy. Clin Neurophysiol. 2024;161:198–210.

118. Ziaeemehr A, Woodman M, Domide L, Petkoski S, Jirsa V, Hashemi M. Virtual Brain Inference (VBI), a flexible and integrative toolkit for efficient probabilistic inference on whole-brain models. Elife. 2025;14:RP106194.

119. Kini LG, Bernabei JM, Mikhail F, et al. Virtual resection predicts surgical outcome for drug-resistant epilepsy. Brain. 2019;142(12):3892–3905.

120. Nissen IA, Millán AP, Stam CJ, et al. Optimization of epilepsy surgery through virtual resections on individual structural brain networks. Sci Rep. 2021;11(1):19025.
121. Wang HE, Triebkorn P, Breyton M, et al. Virtual brain twins: from basic neuroscience to clinical use. Natl Sci Rev. 2024;11(5):nwae079.
122. Dollomaja B, Wang HE, Guye M, Makhalova J, Bartolomei F, Jirsa VK. Virtual epilepsy patient cohort: generation and evaluation. PLoS Comput Biol. 2025;21(4):e1012911.
123. Jirsa V, Wang H, Triebkorn P, et al. Personalised virtual brain models in epilepsy. Lancet Neurol. 2023;22(5):443–454.
124. Dallmer-Zerbe I, Jiruska P, Hlinka J. Personalized dynamic network models of the human brain as a future tool for planning and optimizing epilepsy therapy. Epilepsia. 2023;64(9):2221–2238.
125. Jirsa VK, Proix T, Perdikis D, et al. The virtual epileptic patient: individualized whole-brain models of epilepsy spread. Neuroimage. 2017;145:377–388.
126. Wang HE, Woodman M, Triebkorn P, et al. Delineating epileptogenic networks using brain imaging data and personalized modeling in drug-resistant epilepsy. Sci Transl Med. 2023;15(680):eabp8982.
127. Wang HE, Dollomaja B, Triebkorn P, et al. Virtual brain twins for stimulation in epilepsy. Nat Comput Sci. 2025:1–15.
128. Acharya G, Davis KA, Nozari E. Predictive modeling of evoked intracranial EEG response to medial temporal lobe stimulation in patients with epilepsy. Commun Biol. 2024;7(1):1210.
129. Brodbeck V, Spinelli L, Lascano AM, et al. Electrical source imaging for presurgical focus localization in epilepsy patients with normal MRI. Epilepsia. 2010;51(4):583–591.
130. Foged MT, Martens T, Pinborg LH, et al. Diagnostic added value of electrical source imaging in presurgical evaluation of patients with epilepsy: a prospective study. Clin Neurophysiol. 2020;131(1):324–329.
131. Sharma P, Scherg M, Pinborg LH, et al. Ictal and interictal electric source imaging in pre-surgical evaluation: a prospective study. Eur J Neurol. 2018;25(9):1154–1160.

132. Jiang X, Cai Z, Gonsisko C, Worrell GA, He B. Mapping epileptogenic brain using a unified spatial-temporal-spectral source imaging framework. Proc Natl Acad Sci U S A. 2025;122(50):e2510015122.

133. Abdi-Sargezeh B, Shirani S, Valentin A, Alarcon G, Sanei S. EEG-to-EEG: scalp-to-intracranial EEG translation using a combination of variational autoencoder and generative adversarial networks. Sensors (Basel). 2025;25(2):494.

134. Shirani S, Abdi-Sargezeh B, Valentin A, et al. Distributed beamforming for localization of brain seizure sources from intracranial EEG array. In: 2024 32nd European Signal Processing Conference (EUSIPCO). 2024:1117–1121.

135. Ohlerth AK, Valentin A, Vergani F, Ashkan K, Bastiaanse R. The verb and noun test for peri-operative testing (VAN-POP): standardized language tests for navigated transcranial magnetic stimulation and direct electrical stimulation. Acta Neurochir (Wien). 2020;162(2):397–406.

136. Ryvlin P, Rheims S, Hirsch LJ, Sokolov A, Jehi L. Neuromodulation in epilepsy: state-of-the-art approved therapies. Lancet Neurol. 2021;20(12):1038–1047.

137. Valentín A, García Navarrete E, Chelvarajah R, et al. Deep brain stimulation of the centromedian thalamic nucleus for the treatment of generalized and frontal epilepsies. Epilepsia. 2013;54(10):1823–1833.

138. Valentín A, Selway RP, Amarouche M, et al. Intracranial stimulation for children with epilepsy. Eur J Paediatr Neurol. 2017;21(1):223–231.

139. Morrell MJ. Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. Neurology. 2011;77(13):1295–1304.

140. Bergey GK, Morrell MJ, Mizrahi EM, et al. Long-term treatment with responsive brain stimulation in adults with refractory partial seizures. Neurology. 2015;84(8):810–817.

141. Nair DR, Laxer KD, Weber PB, et al. Nine-year prospective efficacy and safety of brain-responsive neurostimulation for focal epilepsy. Neurology. 2020;95(9):e1244–e1256.

142. Rao VR. Personalizing responsive neurostimulation for epilepsy. J Clin Neurophysiol. 2025;42(6):505–512.

143. Geller EB, Skarpaas TL, Gross RE, et al. Brain-responsive neurostimulation in patients with medically intractable mesial temporal lobe epilepsy. Epilepsia. 2017;58(6):994–1004.

144. Fisher R, Salanova V, Witt T, et al. Electrical stimulation of the anterior nucleus of thalamus for treatment of refractory epilepsy. Epilepsia. 2010;51(5):899–908.

145. Gouveia FV, Warsi NM, Suresh H, Matin R, Ibrahim GM. Neurostimulation treatments for epilepsy: deep brain stimulation, responsive neurostimulation and vagus nerve stimulation. Neurotherapeutics. 2024;21(3):e00308.

146. Liang S, Wang L, Shen K, et al. Targeting brain hubs of ictal fast ripple activity to reduce seizures in patients with drug-resistant epilepsy. Sci Transl Med. 2025;17(830):eadq4423.

147. Little S, Pogosyan A, Neal S, et al. Adaptive deep brain stimulation in advanced Parkinson disease. Ann Neurol. 2013;74(3):449–457.

148. Stanslaski S, Summers RLS, Tonder L, et al. Sensing data and methodology from the Adaptive DBS Algorithm for Personalized Therapy in Parkinson's Disease (ADAPT-PD) clinical trial. NPJ Parkinsons Dis. 2024;10(1):174.

149. Neumann WJ, Gilron R, Little S, Tinkhauser G. Adaptive deep brain stimulation: from experimental evidence toward practical implementation. Mov Disord. 2023;38(6):937–948.

150. Ortiz-Guerrero G, Gregg NM. Biomarkers for epilepsy deep brain stimulation. J Clin Neurophysiol. 2025;42(6):486–492.

151. Okoroafor F, Qiang Z, Rosenke S, et al. Neurophysiologic biomarkers of invasive neuromodulation therapy for epilepsy. Neuromodulation. 2026.

152. Guidotti R, Basti A, Pieramico G, et al. When neuromodulation met control theory. J Neural Eng. 2025;22(1):011001.

153. Trevelyan AJ, Marks VS, Graham RT, Denison T, Jackson A, Smith EH. On brain stimulation in epilepsy. Brain. 2025;148(3):746–752.

154. Fröhlich F, McCormick DA. Endogenous electric fields may guide neocortical network activity. Neuron. 2010;67(1):129–143.
155. Baud MO, Bernard C, Frauscher B, Karoly PJ, Rao VR. Timing is everything: expert opinion on researching epilepsy rhythms by the ILAE Task Force on Chronobiology. Epilepsia. 2026.
156. Dan J, Pale U, Amirshahi A, et al. SzCORE: seizure community open-source research evaluation framework for the validation of electroencephalography-based automated seizure detection algorithms. Epilepsia. 2025;66:14–24.

Tables

Table 1. Comparative framework for methods across the epilepsy sensing-to-control pipeline.

Pipeline stage	Temporal scale / latency	Observability / representation	Non-stationarity / drift	Deployment role
Acquisition and datasets	Hours to years; offline context	Raw sensor sampling; modality, geometry, labels	Dataset shift; longitudinal variability; heterogeneity	Bounds feasible inference and generalization
Preprocessing and artifact suppression	Milliseconds to seconds; may require causal operation	Filtered, re-referenced, artifact-suppressed signals	Reference sensitivity; artifact variation; protocol dependence	Defines the observability boundary
Spontaneous events and evoked responses	Milliseconds to hours; event and warning windows	Event markers; risk states; ER/DR summaries	Threshold drift; calibration burden; early-window sensitivity	Provides control-relevant sensing variables
Connectivity analysis	Seconds to days; windowed or aggregated estimates	SC/FC/EC abstractions; graph/pathway summaries	State dependence; leakage; incomplete sampling	Supports mapping, target ranking, and decision support
State estimation and brain models	Seconds to days; recursive inference and planning timescales	Latent states; source estimates; reduced-order/model-derived variables	Model mismatch; non-identifiability; recalibration needs	Provides state summaries for monitoring and policy support
Closed-loop control	Milliseconds to seconds for action; days to months for tuning	Policy-relevant state estimates; bounded action variables	Policy drift; calibration decay; safety-critical failure modes	Implements constrained stimulation under clinician oversight

Note: ER, early response; DR, delayed response; SC, structural connectivity; FC, functional connectivity; EC, effective connectivity. Four deployment-relevant axes are used for cross-stage comparison: temporal scale and latency constraints, observability and representation, vulnerability to non-stationarity and drift, and deployment role.

Table 2. Dataset properties and downstream methodological consequences.

Data property	Downstream methodological consequence	Evaluation bias / deployment failure mode
Modality, geometry, and spatial sampling		
scEEG (broad coverage; artifact-prone; lower spatial specificity)	Continuous detection prioritizes artifact robustness and conservative operating points; cross-subject pipelines benefit from montage-aware preprocessing	Clip-level AUC can appear optimistic; FP/min and alarm stability often become the true deployment bottleneck
iEEG (patient-specific geometry; limited coverage; variable SNR)	Sampling geometry and referencing shape localization and connectivity estimates; sparse sampling limits identifiable nodes and network completeness	Common-source and reference leakage can create spurious hubs; connectivity-based conclusions are bounded by unobserved regions and incomplete sampling
Temporal coverage and validation regime		
Short curated clips (selected segments; weak temporal context)	Encourages high-capacity models and segment-level tuning; thresholds and calibrations often fail to transfer to continuous monitoring	Selection bias and miscalibration; continuous false-alarm burden is underestimated when evaluation is restricted to curated segments
Long-term continuous recordings (hours to months)	Requires drift handling and time-aware validation; enables risk-state modelling and circadian/multidien analyses	Temporal leakage arises if splits ignore time; deployment-relevant metrics should include FP/min, time-in-warning, and longitudinal stability
Cohort, device, and prevalence structure		
Multi-center / multi-device data (site protocols; hardware differences)	Domain shift motivates harmonization, site-aware validation, and explicit monitoring for dataset drift	Apparent gains can vanish under external validation; hidden confounds may dominate reported improvements
Rare-event prevalence / class imbalance	Operating-point selection, cost-sensitive learning, and calibration become central in low-prevalence regimes	Accuracy and AUC can mislead; PPV collapses and FP/min inflates under realistic event base-rates
Stimulation-while-sensing constraints		
Stimulation during sensing (blanking, saturation, residual artifact decay)	Blanking and residual artifacts constrain early-window observability, evoked-response quantification, and perturbational EC estimation	Protocol-dependent bias; early-response features and directed connectivity claims can be distorted by residual stimulation artifacts
Annotation and supervision regime		
Annotation protocol variability (granularity; inter-rater disagreement)	Label policy sets an error floor; modelling uncertainty or label noise can matter as much as model capacity	Apparent gains may reflect annotation policy rather than algorithmic advance; transfer across datasets becomes brittle when supervision schemes differ

Note: scEEG, scalp EEG; iEEG, intracranial EEG; SNR, signal-to-noise ratio; FC, functional connectivity; EC, effective connectivity; AUC, area under the receiver operating characteristic curve; FP/min, false positives per minute; PPV, positive predictive value.

Table 3. Stimulation-artifact suppression method families and deployment-relevant trade-offs.

Family	Best-fit scenarios	Early-window fidelity potential	Real-time suitability	Minimum reporting requirements
Template / reconstruction	Stereotyped artifacts; stable SPES/CCEP protocols; feasible template updating under limited drift	Medium-High	High	Update rule; alignment; early-window residual metric
Blanking / interpolation	Severe saturation or clipping; safety-driven mitigation; simple causal front-end	Low-Medium	Very high	Blanking window (milliseconds); interpolation method; early samples removed
Spectral / adaptive	Harmonic or narrowband artifacts (common in tACS/tES); separable spectral structure	Low-Medium	Medium-High	Filter causality or zero-phase use; bands; group delay or adaptation rule
BSS / projection / subspace	Multi-channel recordings with coherent artifact footprint; approximately stable artifact subspaces	Medium-High	Medium	Channels retained; subspace window; components removed or projected out
Model / state-space	Non-stationary artifacts; protocol variation; explicit dynamics may improve tracking	High	Medium	Model form; identification procedure; early-window fit or residual error
Deep learning	Large and trustworthy training sets; complex artifacts; explicit protection against leakage and early-window distortion	Medium-High	Medium	Training design; label source; leakage controls; early-window fidelity metric

Note: SPES, single-pulse electrical stimulation; CCEP, cortico-cortical evoked potential; tACS, transcranial alternating current stimulation; tES, transcranial electrical stimulation; BSS, blind source separation. 'Early-window fidelity potential' denotes the potential to preserve interpretable signal structure within the first few milliseconds after stimulation, which is critical for early responses and early cortico-cortical evoked potential components. The final column lists the minimum reporting elements needed to interpret early-response or perturbational effective connectivity claims.

Table 4. Comparative task framework for spontaneous-event sensing and SPES-evoked responses.

Task	Output	Time scale / timing constraint	Deployment-aligned evaluation
IED detection	Event timing; optional subtype or morphology label	Millisecond-scale; low-latency triggering	FP/min, operating-point stability, onset jitter, and sensitivity to inter-rater disagreement and label policy
Seizure detection	Ictal state; optional onset estimate	Seconds; detection latency in seconds	Sensitivity versus false alarms/day, detection latency, and longitudinal stability under drift
Seizure prediction	Risk state / hazard estimate	Hours to days; warning horizon	Time-in-warning, false-warning burden, calibration, and prospective leakage-aware validation
SPES-evoked responses (CCEPs)	Evoked-response summaries; optional ER/DR characterization	Milliseconds to hundreds of milliseconds; early-window fidelity	Early-window fidelity, ER/DR detectability, and cross-pulse/cross-session consistency under protocol variation

Note: IED, interictal epileptiform discharge; SPES, single-pulse electrical stimulation; CCEP, cortico-cortical evoked potential; ER, early response; DR, delayed response; FP/min, false positives per minute. Output, timescale, and deployment-relevant evaluation requirements across spontaneous-event sensing and stimulation-evoked response analysis.

Table 5. Comparative framework for connectivity measures and their deployment roles in epilepsy.

Measurement route / evidence source	Connectivity type	Key assumptions / failure modes	Primary deployment role
DTI tractography	SC	Tractography and parcellation bias; SC constrains plausible pathways but does not determine expressed dynamics uniquely	Mapping: feasibility and safety priors; patient-specific anatomical constraints
Resting-state fMRI / PET (contextual)	FC	Strong state and site dependence; preprocessing sensitivity; hemodynamic or metabolic context may not align with electrophysiology	Mapping / prediction: syndrome-level context, stratification priors, and hypothesis generation rather than direct programming variables
Passive EEG/MEG/iEEG	FC	Leakage and reference dependence; topology and hub structure are sensitive to preprocessing, montage, and spatial coverage	Monitoring: phenotyping and state tracking, provided robustness is demonstrated across preprocessing and state
Generative or model-based inference from observed dynamics	EC	Model mis-specification; non-stationarity; inferred directionality depends on SNR and modeling assumptions rather than constituting direct proof of causality	Prediction / programming support: directed propagation hypotheses used cautiously and interpreted with explicit model dependence
SPES/CCEP (perturbational probing)	EC	Protocol-conditioned directed edges; stimulation parameters, response definitions, and preprocessing can reshape edge existence and strength	Mapping / programming: pathway ranking and programming constraints when reproducibility and early-window fidelity are demonstrated
Spike/seizure propagation (temporal precedence)	EC	Detection bias, incomplete sampling, and strong state dependence; inferred directionality reflects recruitment dynamics rather than generic connectivity	Prediction / programming support: outcome-relevant recruitment pathways that complement perturbational EC and structural constraints

Note: DTI, diffusion tensor imaging; SC, structural connectivity; fMRI, functional magnetic resonance imaging; PET, positron emission tomography; FC, functional connectivity; EEG, electroencephalography; MEG, magnetoencephalography; iEEG, intracranial EEG; EC, effective connectivity; SNR, signal-to-noise ratio; SPES, single-pulse electrical stimulation; CCEP, cortico-cortical evoked potential. Connectivity type, dominant interpretability limits, and primary role in mapping, monitoring, prediction, or programming for each measurement route or evidence source.

Table 6. Comparative framework for brain models and cross-modal transformations in epilepsy.

Model family	Dynamical commitment	Observability / representation	Dominant failure modes	Primary deployment role
Latent-state models	Phenomenological	Sensor-level to low-dimensional latent state	State mislabeling, calibration drift, non-stationarity, limited interpretability	Monitoring / triggering / risk estimation
Reduced-order or surrogate control models	Intermediate	Low-dimensional state or reduced dynamics derived from observed activity	Local validity only, state drift, limited generalization outside fitted regime	Policy design / constrained control prototyping
Mechanistic seizure models	Mechanistic	Latent physiological variables and structured seizure dynamics	Non-identifiability, parameter uncertainty, mismatch to routine clinical observables	Explanation / hypothesis testing / planning
Network ictogenicity / virtual resection	Intermediate to mechanistic	Patient-specific network propensity measures	Connectome bias, incomplete sampling, model dependence, limited external validation	Planning / candidate target ranking
Electrical source imaging and source-space transformations	Transformational	Sensor-to-source mapping	Head-model dependence, preprocessing sensitivity, inverse ill-posedness, event-selection bias	Mapping / localization support
Cross-modal or generative transformations	Transformational / data-driven	Sensor-to-source or sensor-to-intracranial representations	Domain shift, weak ground truth, leakage, visually plausible but decision-irrelevant outputs	Exploratory mapping / representation alignment
Personalized virtual-brain frameworks	Mechanistic / integrative	Virtual patient-specific latent and network states	Uncertainty propagation, parameter non-uniqueness, incomplete sampling, state mismatch	In silico testing / scenario exploration / pre-intervention planning

Note: Deployment relevance depends on observability, identifiability, robustness under drift, and supported decision role, rather than on biological detail alone.

Table 7. Policy interface for closed-loop neuromodulation.

Policy component	Deployment-oriented definition
Inputs	Sensed quantities available to the policy: event triggers (IED or seizure detections), slower context or risk covariates (e.g., circadian or multidien phase, hazard), and network-derived priors or constraints (e.g., perturbational or propagation-derived EC summaries).
State estimation	Low-dimensional inferred state used for control: biomarker vector or discrete latent mode, together with calibration, uncertainty quantification, and drift monitoring suitable for longitudinal use.
Action space	Bounded stimulation actions specifying when (timing or scheduling), where (contact or target), and how (amplitude, pulse width, frequency, burst structure) stimulation may be delivered, including optional mode switching under explicit rules.
Constraints	Hard limits and supervisory bounds: charge density, duty cycle, refractory or cool-down rules, power and latency budgets, and clinician-defined restrictions or overrides.
Fallback behavior	Conservative default behavior under uncertainty, signal loss, or instability: safe-mode actions, stimulation suppression, and explicit criteria for recovery or re-entry into adaptive operation.
Learning and update	Governed adaptation mechanisms: clinician-in-the-loop reprogramming or, where justified, bounded autonomous adaptation, supported by auditability, logging, versioning, and rollback.

Note: IED, interictal epileptiform discharge; EC, effective connectivity. Closed-loop systems are treated as governed control policies defined by sensed inputs, inferred states, bounded actions, safety constraints, fallback logic, and explicit update mechanisms, rather than simply as detectors linked to stimulators.

Figure

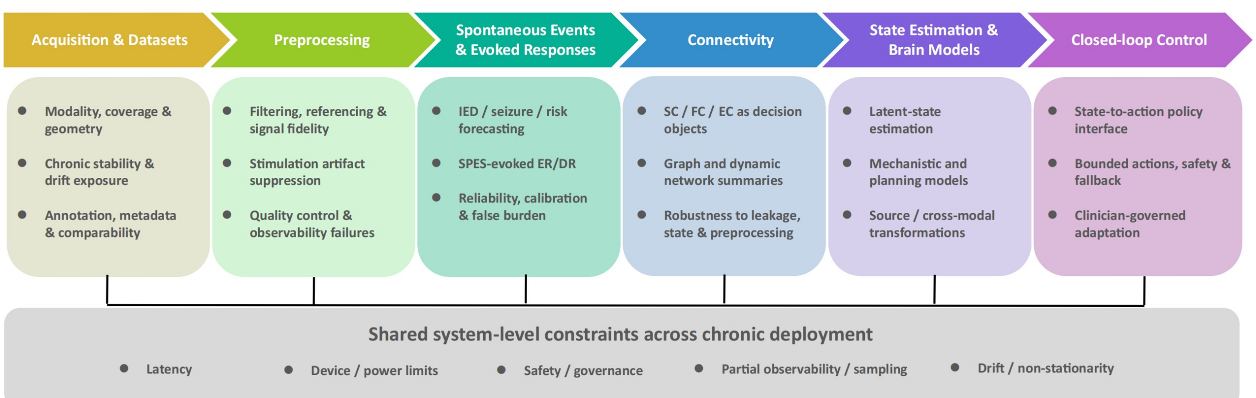

Figure 1. Deployment-oriented sensing-to-control pipeline in epilepsy, spanning acquisition, preprocessing, sensing, connectivity, brain modeling, and closed-loop control under shared constraints of latency, observability, safety, and drift.