From brain signals to smart modulation: A scoping review of electroencephalography 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, *], Ioannis Stavropoulos[1, 5, *]

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

[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, London, UK

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

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

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

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

[7]School of Biomedical Engineering, Tsinghua University, Beijing, China

***Correspondence to**: Ioannis Stavropoulos, MD, MSc, PhD, ioannis.stavropoulos@kcl.ac.uk; ORCID: 0000-0002-6920-1826

Abstract

Despite substantial recent progress in epilepsy related technologies and algorithms, strong retrospective performance often fails to translate into reliable long term clinical use 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 translational 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. These are distinct from the broader research agenda summarized by the four open questions below. 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.

Keywords: epilepsy; neuromodulation; neurostimulation; closed-loop control; seizure detection; seizure forecasting; electrophysiological biomarkers; connectivity analysis; brain modeling; clinical deployment

Open questions

1. How can pipeline-induced uncertainty be evaluated and reported across epilepsy sensing-to-control pipelines? A central unresolved issue is how constraints introduced by recording, preprocessing, artifact suppression, and event sensing propagate into downstream localization, modeling, stimulation programming, and closed-loop control.

2. What constitutes a reliable control variable under partial observability and non-stationarity? Candidate control variables may be derived from spontaneous events, stimulation-evoked responses, connectivity measures, or learned latent states, but their stability, interpretability, and clinical validity may change across brain states, patients, recording contexts, and stimulation history.

3. How can clinically trustworthy operating points be maintained during chronic use? Long term systems must detect and respond to degradation in calibration, false alarm burden, biomarker reliability, and uncertainty when seizures are rare, annotations are incomplete, and patient-specific dynamics evolve over time.

4. How much bounded autonomy can adaptive neuromodulation safely assume? Future systems will need to define which updates can occur automatically, which require clinician approval, and how escalation, rollback, and safety constraints should operate when thresholds, policies, or state definitions change during chronic use.

1 Introduction

Epilepsy affects over 50 million people worldwide and remains refractory to pharmacological treatment in approximately a third of patients. [1;2] 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. These methods should accordingly be treated as parts of a sensing-to-control 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.[3–6] 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.

Across these stages, deployment context introduces a further layer of constraint that cuts across modalities and algorithmic tasks. Three recurring regimes are considered throughout this review. For in-hospital or epilepsy monitoring unit settings, stable external power supply, dense recording infrastructure, and continuous clinician oversight are usually available, so the dominant constraints are sensitivity, interpretability, and short-horizon decision support. In home, ambulatory, or wearable settings, prolonged monitoring, changing impedance, environmental noise, adherence gaps, and intermittent computation shift the priority toward false alarm burden, calibration stability, and robustness to drift. In fully implantable closed-loop systems, strict limits on power, memory, charge delivery, sensing quality, and real-time latency make bounded autonomy, safety constraints, and hardware-algorithm co-design central. Therefore, the constraints discussed in later sections are mapped back to these regimes, because the same algorithmic choice may carry different evidential and practical implications across deployment contexts.

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

This review adopts a structured scoping approach and an engineering-oriented conceptual synthesis to organize key methods within the epilepsy sensing-to-control pipeline (Figure 1 and Table 1).

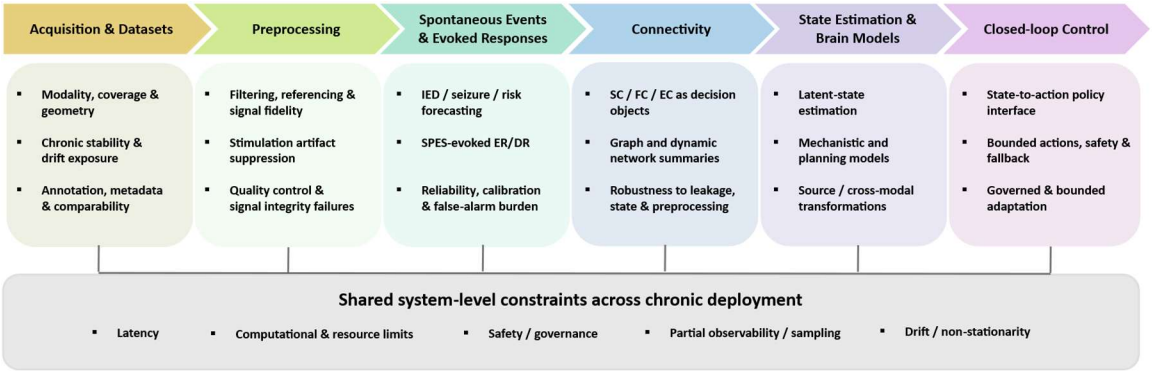

Figure 1. Sensing-to-control pipeline for algorithm-guided neuromodulation in epilepsy. End-to-end pipeline from acquisition to closed-loop control, operating under shared chronic-deployment constraints: latency, computational and resource limits, safety/governance, partial observability/sampling limitations, and drift/non-stationarity.

2.1 Literature search and selection criteria

We conducted a structured scoping search across PubMed, Institute of Electrical and Electronics Engineers Xplore (IEEE Xplore), ScienceDirect, and Google Scholar. The search covered publications from January 1990 to June 2, 2026, including ahead-of-print and in-press articles available by that date. Search terms were organized into three broad blocks: (1) epilepsy and neuromodulation terms, including epilepsy, seizure, epileptiform activity, neuromodulation, stimulation, responsive neuromodulation (RNS), deep brain stimulation (DBS), vagus nerve stimulation (VNS), and closed-loop stimulation; (2) recording and sensing modalities, including scalp electroencephalography (scEEG), intracranial electroencephalography (iEEG), electrocorticography (ECoG), stereo-electroencephalography (SEEG), local field potentials (LFPs), chronic electroencephalography (chronic EEG), ambulatory electroencephalography (ambulatory EEG), and stimulation-while-sensing recordings; and (3) methodological and deployment targets, including artifact handling, event detection, seizure detection, seizure forecasting, evoked-response quantification, single-

pulse electrical stimulation (SPES), SPES-evoked responses, and cortico-cortical evoked potential (CCEP), source localization, connectivity analysis, brain modeling, machine learning (ML), deep learning (DL), adaptive control, calibration, drift, and clinical deployment. The search strings for each database are provided in Supplementary Table S1.

Studies were considered eligible when they addressed epilepsy related electrophysiological sensing, inference, modeling, or neuromodulation and were relevant to at least one stage of the sensing-to-control pipeline. Non-epilepsy studies were included only when they provided a directly relevant methodological template, such as adaptive deep brain stimulation as an example of governed biomarker-based control. Eligible references were then classified as either primary deployment-facing evidence or contextual evidence. Primary deployment-facing evidence was defined as original epilepsy related data, device evaluation, or model evaluation that met at least one predefined criterion: continuous or long-duration data, temporally separated, chronological, or prospective validation, explicit operating-point selection, latency or safety reporting, stimulation-while-sensing evaluation, longitudinal follow-up, or robustness testing under non-stationarity, cross-session variation, or calibration decay. Contextual evidence included reviews, consensus papers, foundational methodological studies, mechanistic studies, dataset or resource papers, and reporting frameworks used to define method families, clarify terminology, or support conceptual rationale. Contextual evidence was not used as the primary basis for claims about deployment robustness or clinical reliability.

Database searches identified 5,293 records; after deduplication, 4,052 unique records remained. Title screening retained 1,522 records for abstract review, of which 328 were assessed in full text and 205 were included in the final scoping analysis. The selection pathway is summarized in a PRISMA-ScR flow diagram (Preferred Reporting Items for Systematic Reviews and Meta-Analyses Extension for Scoping Reviews; Figure 2). Included references were classified by evidence type, and deployment-facing studies

were further classified by validation level, dominant recording modality, and study design (Table 2).

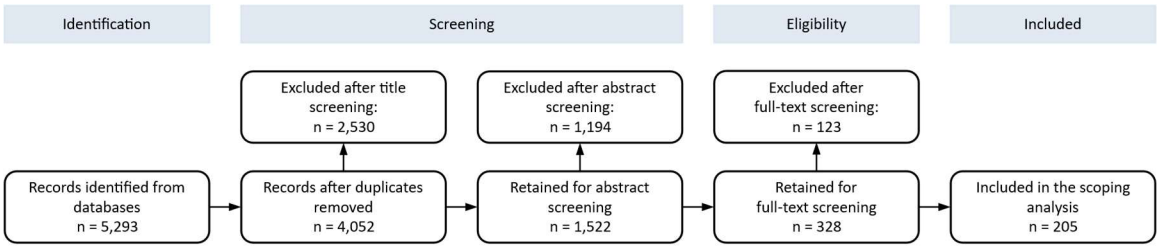

Figure 2. A PRISMA-ScR flow diagram of the selection process.

2.2 Operational four axis framework

To compare heterogeneous algorithms across the sensing-to-control pipeline, this review organizes methods along four operational axes (temporal scale and latency, observability and representation, vulnerability to non-stationarity and drift, and deployment role), each operationalized as a set of reportable descriptors.

The first axis, temporal scale and latency constraints, describes both the physiological time scale of the target variable and the allowable delay between sensing, inference, and action. Relevant descriptors include the analysis-window length, whether processing is causal or non-causal, processing latency, warning horizon, stimulation-trigger delay, and update interval. The second axis, observability and representation, describes which variables are directly measured and which are inferred from limited sensors, spatial sampling, labels, or model assumptions. This distinction is important because clinically relevant epileptic-network states are typically inferred from finite sensors, incomplete anatomical coverage, noisy measurements, imperfect labels, and model-dependent representations rather than measured directly (see Box 1). Relevant descriptors include recording modality, sensor density and anatomical coverage, SNR, source ambiguity, feature or latent-state definition, label availability, and uncertainty reporting. The third axis, vulnerability to non-stationarity and drift, describes whether the input distribution, biomarker expression, model output, or stimulation-response relationship changes across state, session, treatment, or chronic use. Relevant descriptors include chronological validation gap, cross-session degradation, recalibration

frequency, threshold drift, false alarm drift, and robustness to artifacts, sleep, medication, circadian or multidien rhythms, and stimulation effects. The fourth axis, deployment role, distinguishes offline mapping, clinician decision support, continuous monitoring, prospective risk estimation, real-time closed-loop control, and adaptive stimulation. Relevant descriptors include online feasibility, clinician-in-the-loop requirements, safety envelope, override or rollback mechanisms, and the validation level required to support the intended claim.

These descriptors are summarized across pipeline stages in Table 1. To illustrate how the same framework can be applied to concrete method families, Table 3 provides worked examples spanning retrospective scEEG seizure detection, CCEP mapping, and implantable responsive stimulation. Because reportable descriptors do not by themselves establish the evidential strength of a deployment claim, Table 4 further distinguishes retrospective discrimination, quasi-online testing, prospective online operation, and chronic clinical deployment, clarifying what each level can and cannot support. [7–10]

Box 1. Engineering interpretation of partial observability

Definition: A latent epileptic-network state is observed only through a limited montage. Partial observability denotes the mismatch between the clinically relevant state and the sensor-derived measurement.[11;12]

Reportable surrogates: Useful proxies include spatial coverage ratio, captured-event fraction from concurrent scalp-intracranial recordings,[13;14] scalp source ambiguity or leakage,[15;16] and biomarker stability and calibration.[11;17]

Modality trade-off: scEEG offers broad but low-specificity coverage; iEEG offers high local specificity but sparse, implantation-dependent sampling. Neither provides full network observability.

3 Electrophysiological Data Acquisition and Datasets

Electrophysiological data define the empirical basis on which epilepsy algorithms are developed, evaluated, and interpreted. Recording modality, spatial sampling, temporal

coverage, annotation quality, and dataset structure jointly determine which claims can be tested and how far those claims can be generalized. However, in epilepsy research, data are not only heterogeneous across studies in recording conditions, label quality, and temporal structure but are also unevenly available across tasks and deployment contexts. Large retrospective EEG resources are widely used in seizure detection and prediction research, whereas chronic implantable recordings are represented by small, specialized cohorts, and stimulation-while-sensing data are available through a small number of public resources. [24–27] As a result, benchmarking often remains shaped more by data availability than by clinical need.

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 iEEG, the latter encompassing ECoG and SEEG. These modalities differ not only in invasiveness, spatial resolution and coverage, signal-to-noise (SNR), and feasible recording duration, but also in the kinds of inference they can support. For clinical use, they provide complementary rather than interchangeable forms of evidence.

iEEG 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.[28] 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. [29] 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 large-scale data collection because it is non-invasive. This has made it the dominant substrate for population-scale datasets and many benchmarking efforts. [24] 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. [13;30] 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.[31–37] These properties matter directly for forecasting, calibration, and adaptive neuromodulation because they determine which failure modes become visible during evaluation. Therefore, chronic ambulatory recordings are especially important for assessing long term forecasting realism, calibration stability, and real-world monitoring burden, whereas chronic implantable-device recordings are more directly relevant to sensing substrates used for adaptive neuromodulation and closed-loop control. Intracranial stimulation-while-sensing datasets are similarly important for quantifying SPES-evoked responses, supporting cortico-cortical mapping, and estimating stimulation-derived effective connectivity. [20;26;27;38] Across these acquisition contexts, 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 and clinical-use claims they can support. Their utility depends on several interacting factors, including recording duration, sampling frequency, recording modality, annotation precision, temporal continuity, and the context in which data were collected. Table 5 summarizes representative resources and their clinical-use characteristics.

Large clinical scEEG datasets remain among the most widely used public resources in epilepsy. The Temple University Hospital EEG Corpus and its seizure-annotated subsets have become standard benchmarks because they expose algorithms to substantial variability in recording conditions and patient state. [39;40] Open infrastructures such as PhysioNet (https://physionet.org/) 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. [18] More recently, structured clinical EEG resources, such as the Harvard Electroencephalography Database (https://github.com/bdsp-core/Harvard-EEG-Database-Tools), have added richer clinical metadata to this landscape. [41] These short term or clinically sampled retrospective resources support benchmarking, model comparison, and large-scale offline analysis, but their temporal structure and label granularity limit what can be inferred about long term reliability.

Longer duration resources are especially important for seizure prediction and forecasting because they preserve temporal context for time-aware validation and circadian or multidien analyses.[24;42] EPILEPSIAE is a prominent long term presurgical monitoring resource, while NeuroVista remains notable as one of the few chronic invasive recording datasets used for seizure forecasting over months to years. [25;42;43] Their scarcity, selected cohorts, and limited accessibility nevertheless constrain how confidently long term forecasting robustness can be generalized across broader epilepsy populations. [24]

Other public resources are organized around more specific analytical 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 scEEG dataset introduced by Lin et al., particularly given the relative scarcity of public EEG datasets with expert IED annotations and spatial information. [44] A complementary intracranial resource is provided by Falach et al., who released annotated iEEG sleep recordings from two medical centers with sub-second, multi-channel IED annotations in hippocampal and medial temporal lobe regions, together with inter-rater reliability assessment. [45] A more recent high-frequency oscillation (HFO)-focused benchmark resource is Omni-iEEG, which harmonizes presurgical iEEG recordings from 302 patients across eight leading epilepsy centers and provides over 36K expert-validated pathological event annotations, focusing on HFOs co-occurring with spikes (spkHFOs). Unlike these IED-centered datasets, Omni-iEEG defines benchmark tasks for pathological event classification and brain-region identification, enabling systematic, multicenter evaluation of both biomarker-driven and end-to-end data-driven approaches.[46] Others support sensing-while-stimulating paradigms, where stimulation-evoked responses are used to localize seizure onset zones and characterize network behavior under controlled perturbation. In epilepsy, SPES-based datasets are one of the clearer examples of this kind of resource, with PRIOS providing protocol-specific iEEG Brain Imaging Data Structure (iEEGBIDS) CCEP recordings and Parmigiani et al. providing simultaneous SEEG and high-density scEEG recordings that serve as a cross-scale reference for assessing how stimulation parameters shape CCEP responses. [26;27] Taken together, these resources support more precise but modality- and protocol-bounded questions, although the same protocol specificity limits generalizability beyond the target paradigm.

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. [47] The iEEG extension of the BIDS provides a formal structure for organizing intracranial

electrophysiology data and supports interoperability, scalable preprocessing, and cross-study reuse. [48] These infrastructures reduce fragmentation and support reproducible reuse, but the evidential value of a dataset still depends on its annotation structure, cohort composition, and validation context, as discussed below.

3.3 Engineering consequences

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 shape threshold selection, the transferability of calibration across sessions or devices, and the kinds of drift or clinical-use failure modes that can be assessed during validation. [7;8;24] Table 6 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 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.

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. [49;50]

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.[50;51] Therefore, filter type, cutoff definition, order or length, transition bandwidth, phase properties, causality, and filtering direction should be reported explicitly, especially when latency or waveform morphology is interpreted. [51] This issue is particularly important for high-frequency biomarkers, because sharp transients or non-sinusoidal waveforms can be transformed by high-pass or band-pass filtering into ripple-like events, creating false-positive oscillatory features. [52] Even modest choices, such as a highpass cutoff of 0.1 versus 1 Hz, can reshape slow-wave morphology, shift feature distributions, and modify connectivity estimates.[51]

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. [49] In ECoG and SEEG, bipolar or other local references are often preferred to reduce spatial mixing and hardware-related contamination. [53] 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.[49] Iterative rejection and rereferencing can improve

stability. [53] 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 structures, 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, and increasingly learning-based denoising approaches.[30;49;50;53] Independent component analysis (ICA) used for blind source separation remains widely used, including for component rejection in epileptic EEG and supervised denoisers have been trained to approximate ICA-cleaned signals. [54;55] 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. [20;56] Their morphology depends on stimulation amplitude, pulse width, phase balance, electrode geometry, impedance, and electrode-tissue interface properties, and they can obscure neural activity in the first milliseconds after stimulation.[20;56] This is precisely the window required for quantifying early components of SPES-evoked responses and for estimating perturbational effective connectivity.[20;38;57] Similar observability constraints also arise in DBS sensing and other invasive stimulation settings.[58] Related, though not identical, problems also appear in transcranial electrical stimulation, especially transcranial alternating current stimulation (tACS), and in transcranial magnetic stimulation (TMS) EEG. [59–61]

Hardware-level mitigation can reduce artifact burden at acquisition, but it does not guarantee recovery of the earliest usable post-stimulus signal. Its effectiveness remains constrained by safety, engineering complexity, and protocol dependence. [56;58] 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 7).

4.2.1 Conservative suppression

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. [61] Template methods can be effective when artifacts are sufficiently stereotyped, but they lose reliability when waveform shape drifts across pulses, sessions, or impedance states. For online use, blanking and interpolation are computationally light and can be implemented when artifact timing is known, but they discard or reconstruct the contaminated interval, offering only limited assurance for early-window signal recovery. Simple template subtraction can be low-latency, but remains vulnerable to artifact-shape drift and template-update errors. [56;58;61]

4.2.2 Reconstruction and subspace-based suppression

When early-window fidelity matters, more elaborate approaches may be justified. Biophysical reconstruction has been used to recover early components of SPES-evoked responses, and sparse approaches, such as matching-pursuit-based artifact reconstruction and removal, attempt to separate artifact, neural signal, and noise with high temporal fidelity. [20;62] 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. [61;63] Parametric and state-space formulations model non-stationary artifact dynamics explicitly and may improve robustness under protocol variation. [60;64] 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. Their

computational burden is usually higher than conservative suppression and depends on model complexity, but they are often the most relevant class when residual artifact amplitude and retention of early neural responses need to be balanced explicitly.

4.2.3 Learning-based suppression

Learning-based suppression methods can capture nonlinear and modality-dependent artifact structure.[64] However, they also introduce distinct risks, including over-smoothing, response-shape distortion, leakage between training and evaluation conditions, and limited generalization beyond controlled or semi-synthetic settings. For online or closed-loop use, inference latency, model size, memory requirements, power consumption, and cross-session generalization should be evaluated alongside residual artifact suppression and preservation of physiologically relevant temporal and spectral response structure. [56;58;64]

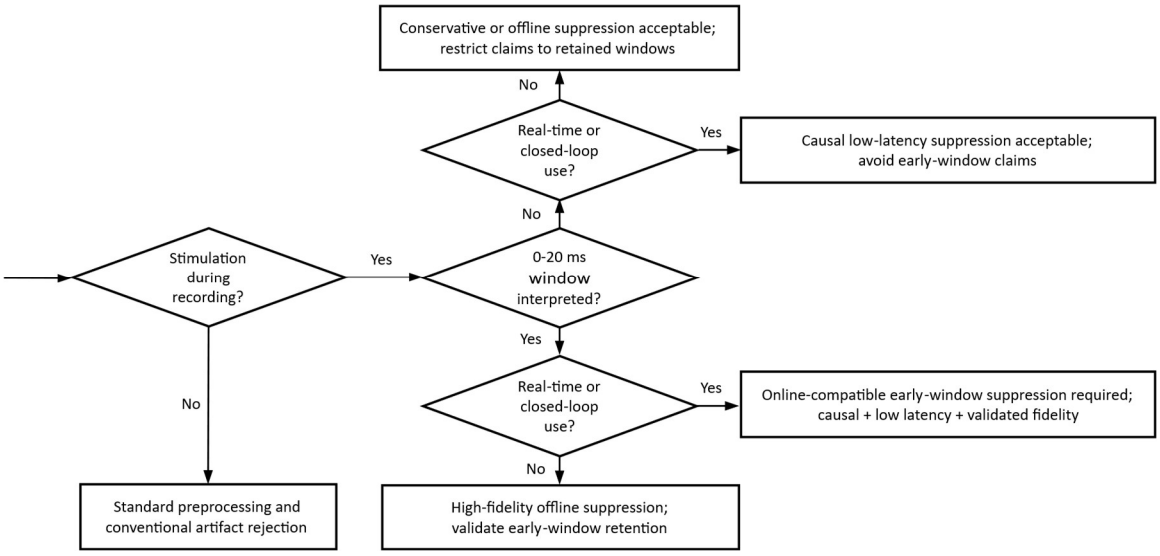

Figure 3. Decision aid for selecting stimulation-artifact suppression strategies. Early-window interpretation defines the primary claim boundary, while real-time or closed-loop use adds constraints on causality, latency, and online feasibility. When both apply, requirements are cumulative rather than alternative. Table 7 gives method-family comparisons and reporting details.

4.3 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 the primary endpoint, as in quantification of early responses 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.[20;61]

Beyond task requirements, deployment context further constrains which suppression strategies are suitable. In in-hospital monitoring, interpretability, reproducibility, and compatibility with clinical workflows may be more important than maximal artifact recovery. In long term home or wearable recordings, robustness to changing impedance, movement, environmental noise, and cross-session drift become central. In fully implantable closed-loop systems, artifact suppression is not only a post-processing problem but also a front-end hardware and system-design constraint, involving amplifier recovery, common-mode and differential-mode artifact rejection, power consumption, memory use, and real-time latency. [56;58;65;66] These setting-specific constraints are summarized in Figure 3, which links stimulation-recording configuration, deployment context, and early-window interpretation requirements to appropriate suppression and reporting choices.

4.4 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.[20] This requirement is particularly important for analyses

of early SPES-evoked responses in the 0-20 ms post-stimulation window, where the measured response is closest to the stimulation artifact and is most sensitive to blanking, interpolation, filtering, and saturation recovery. Stimulation artifact amplitude itself can carry distance-dependent information and may contribute to measured early-response amplitude through volume conduction. [67] Accordingly, studies that interpret this early window should report artifact-handling choices at a level sufficient to judge whether the measured signal remains physiologically interpretable; method-specific reporting elements are summarized in Table 7. This reporting requirement is reinforced by recent CCEP methodology studies showing that preprocessing choices and stimulation-artifact handling vary across studies and can materially affect CCEP quantification.[21;68]

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 online-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 validation-level evaluation (Table 8). 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 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 SNR in scEEG, inter-rater disagreement, and cross-site or cross-device shift. [69] 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.[13]

Seizure detection operates on longer time scales and must remain reliable under behavioral variability, changing vigilance state, and months-long non-stationarity. [70;71] 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.[72–78] 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.[79–81]

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 multi-channel structure.[19;82–87] However, the broader methodological overlap does not remove the fact that the three tasks differ in output structure, time horizon, and deployment burden. Recent patient-specific studies further suggest that chronological retraining, concept-drift adaptation, and incremental updating are part of the method design space for event-sensing models, especially when signals and background states change over time. [71;77;78]

5.1.2 Evaluation and deployment implications

For IED detection, the relevant output is a time-stamped event, so low latency and onset jitter matter more than retrospective discrimination alone, with task-specific reporting expectations summarized in Table 8. In long continuous recordings, the main burden is the false positive rate during long continuous recordings, because clinician review scales directly with false alarm frequency. Accordingly, the most relevant evidence comes from long continuous streams interpreted at an explicit validation level, as summarized in Table 4.

For seizure detection, the central trade-off is detection latency versus false alarms over long continuous streams. The main operational challenge is maintaining that trade-off as background state, medication, vigilance, and signal characteristics change over time. Continuous stream evaluation should therefore be interpreted using the validation levels summarized in Table 4 rather than relying on strong performance on isolated segments. For seizure detection in long term use, incremental-learning approaches are relevant because patient-specific updating can address distribution shift, while replay-based mechanisms aim to preserve rare seizure related patterns and reduce catastrophic forgetting during continual model updates.[71] Ultra-long-term subcutaneous EEG further illustrates this review burden: automated seizure detection is needed to make months-long recordings clinically reviewable, but recent deep learning

evidence remains based on retrospective offline validation, so latency, real-time feasibility, and prospective clinical impact still require separate evaluation. [88]

For seizure prediction, the relevant output is not a discrete event but a risk estimate or warning state. Therefore, its evaluation should be aligned with the validation-level requirements summarized in Table 4, together with the task-specific framework in Table 8. Circadian and multidien structure may provide one relatively stable basis for probabilistic risk estimation in long term settings, making risk forecasting more clinically plausible than brittle binary prediction. [25;43;74;89–93] Chronic ambulatory iEEG has also supported more specific individualized risk tasks, including seizure-cluster and next-seizure prediction, although retrospective cross-validation limits such evidence to methodological support rather than prospective validation.[94] Prospective mobile and wearable forecasting based on cycles in seizure occurrence and heart rate provides one of the few non-invasive examples evaluated on blinded future data, whereas seizure-day risk forecasting from sleep features measured with a smart shirt represents a more controlled route within an epilepsy monitoring unit (EMU), approximating prospective wearable evaluation. [93;95] The principal methodological risk is overestimating prospective utility from retrospective structure, particularly when temporal leakage, post-hoc pre-ictal windows, unstable definitions of seizure risk, or weak operating-point calibration inflate apparent predictability. [9;10;76;90;96] Accordingly, the corresponding evidential standard follows the prospective- and quasi-prospective validation requirements summarized above.[97;98]

Taken together, task-specific reporting should be interpreted through Table 8 and the validation levels in Table 4, which distinguish task outputs from validation-level requirements. Without this mapping, 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. SPES-evoked responses broadly refer to neural activity elicited

by single-pulse electrical stimulation across intracranial stimulation and recording sites. CCEPs are a more specific subset in which stimulation and response sites are interpreted as cortico-cortical interactions for effective-connectivity mapping. Within SPES-evoked responses, early responses (ERs) and delayed responses (DRs) denote shorter- and longer-latency response classes, respectively. ERs generally refer to earlier, relatively fixed-latency components often associated with more direct pathways, whereas DRs refer to later and more variable responses that may reflect broader network reverberation and can resemble spontaneous IED-like activity in epileptogenic tissue.[99–101] However, in practice, the ER/DR distinction is not always clear-cut, so conservative operational definitions are important when SPES-evoked responses are quantified or compared across studies. SPES-evoked responses can be used to probe functional organization, epileptogenicity, and directed network effects. [102–106] A prospective study has linked abnormal SPES-evoked responses to epileptogenic or structurally abnormal cortex and, in some settings, to postsurgical outcome, with similar localization value reported in a pediatric cohort.[107;108]

When SPES-evoked responses are used to support neuromodulation decisions, they must be represented robustly enough to support component classification, mapping, and downstream connectivity analysis. 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. [102;103;106] More structured representations have also been proposed. Basis profile curves and canonical response parameterization, for example, aim to capture response shape using combinations of canonical waveform components. [109;110] Transfer-function and response-shape approaches extend this further when evoked responses are used to construct downstream network objects.[111;112] 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. [113] 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 ERs from DRs within SPES-evoked responses or to construct downstream connectivity measures. Machine learning approaches may support localization, response typing, and automated response detection or standardization.[112;114–117] However, they do not remove the requirement that the analyzed response window itself be sufficiently reliable for interpretation.

6 Connectivity Analysis and Networks of Epilepsy

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. This makes connectivity 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, periictal state, and longer term drift. These constraints place strong limits on how connectivity measures can be interpreted and used.

We first distinguish structural connectivity (SC), functional connectivity (FC), and effective connectivity (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 EC in decision support (Table 9).

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

SC, FC, and EC answer different questions and are not interchangeable. 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. 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. [118]

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. [119] 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. [120] Anatomical plausibility is an important constraint, but not a ground-truth oracle for electrophysiological connectivity. [119;120]

A further boundary is methodological rather than biological. Zero-lag correlations may arise from volume conduction or source leakage rather than true interaction. These failure modes vary systematically across estimator families. Correlation and coherence are sensitive to common-reference contamination, volume conduction, SNR differences, and shared or unobserved input, and can therefore reflect spurious coupling rather than direct interaction. Phase-synchrony measures such as the phase-locking value reduce amplitude dependence but remain vulnerable to common-reference effects, volume conduction, noise, and sample-size bias. Imaginary coherence, the phase lag index, and the weighted phase lag index reduce zero-lag confounds, but their conservative design can miss genuine interactions with near-zero phase delay or interactions not proximal to the sampled sensors. Directed measures such as Granger-causality-based approaches and transfer entropy add further dependencies on model assumptions, model order or estimator choice, SNR, data quality, and unobserved common inputs; therefore, inferred directionality should be interpreted as model-conditioned evidence rather than direct causal proof. Graph-theoretic summaries inherit the assumptions of the upstream connectivity estimator and add sensitivity to node definition, parcellation, thresholding, network density, and spatial sampling; therefore, node centrality and hub rankings should be treated as pipeline-dependent outputs rather than intrinsic network properties.[15;121–126]

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.[16] Referencing adds a further layer of sensitivity. Common-average, bipolar,

and Laplacian montages can change edge weights, graph structure, and apparent hubs. [127] This sensitivity is not only theoretical. In a large-scale iEEG study involving 125 patients and 48 preprocessing pipelines, rereferencing and connectivity choices substantially altered functional network estimates, robustness to incomplete spatial sampling, and susceptibility to spurious correlations, and different pipelines varied markedly in their ability to lateralize the clinically defined seizure-onset hemisphere in temporal lobe epilepsy. [126] In iEEG, these issues are compounded by limited coverage and sampling geometry. Incomplete coverage can introduce missing-node bias, while dense sampling of clinically suspected regions can create apparent hubs that partly reflect implantation strategy rather than true network centrality. Transparent reporting of coverage, distance-control analyses, and montage robustness is particularly important.

When connectivity estimates are used to support neuromodulation decisions, they should demonstrate robustness to plausible referencing schemes, leakage-control choices, thresholding or graph-density variation, and physiological state. These requirements differ across recording modalities. Scalp EEG provides non-invasive and relatively broad coverage, but network estimates remain limited by volume conduction, source mixing, inverse-model dependence, and lower effective spatial specificity than intracranial recordings. [15;121;122] Accordingly, scalp EEG networks are most appropriate for offline or slow-timescale monitoring and hypothesis generation. By contrast, iEEG provides higher local SNR and spatial specificity, but its sparse, hypothesis-driven coverage and reference sensitivity make inferred topology conditional on implantation geometry and preprocessing choices. Electrode density can further alter the clinically inferred seizure onset zone (SOZ) extent and inter-rater agreement, so coverage and density should be reported when iEEG network summaries are used for patient-specific mapping, candidate contact prioritization, or periodic programming review. [128]

6.2 Graph representations and dynamic network summaries

These coverage and preprocessing constraints carry over to the graph and dynamic summaries considered next. Graph analysis converts a connectivity matrix into a

network representation in which brain regions or parcels are treated as nodes, and estimated structural or functional interactions are treated as edges. [129] This representation allows the same connectivity data to be summarized at multiple levels: nodal centrality quantifies how central or relevant a region is within the network, modular structure describes how strongly subnetworks are segregated, and global measures such as path length characterize overall communication efficiency across the graph.[129] In epilepsy, this framing is particularly useful because hub regions, identified as highly central nodes and sometimes organized as densely interconnected rich-club structures, have been implicated in seizure spread, cognitive dysfunction, and postsurgical seizure outcome.[129;130]

Given the estimator- and pipeline-dependence established above, graph measures inherit further sensitivity to node definition, parcellation, thresholding, and graph-density choices.[124;125] Their clinical value is strongest when they are tied to specific decision problems, such as ranking candidate seizure-onset regions or predicting whether resection of highly connected nodes is associated with outcome. [131;132] For example, the connectivity epileptogenicity index provides a ranked estimate of candidate seizure-onset regions from SEEG recordings, with better agreement with visually defined SOZ than the original epileptogenicity index for slow-onset seizure patterns. [131] CCEP-based directed graphs may also support localization hypotheses, but their interpretation remains dependent on stimulation protocol, parcellation, sampling coverage, and preprocessing choices. [112;126]

A further limitation is that many graph summaries are derived from connectivity estimates that are averaged over time or interpreted as if they reflected a relatively stable network configuration. In practice, statistical coupling can vary with vigilance state, peri-ictal context, antiepileptic-drug tapering or other treatment related effects, and longer seizure-risk cycles.[119;133] Dynamic FC methods address this limitation by asking whether a fixed set of nodes passes through recurrent connectivity or activity states over time, and by quantifying how long these states persist and how often transitions occur. [134;135] Within a single recording session, sliding-window clustering and

hidden Markov models (HMMs) have been used to identify such states and describe them using summaries such as dwell time, fractional occupancy, and transition probability. [134;135] At a slower monitoring scale, daily vigilance-controlled iEEG connectivity has been used to distinguish interictal from preictal days and generate patient-specific probabilistic seizure-risk forecasts. [133] These dynamic summaries can reveal how network patterns vary over time. For clinical use, the key question is whether the identified states are reproducible in independent recordings or prospective data, can be linked to plausible physiological variation rather than analysis choices, and are described with enough methodological detail to allow replication.

6.3 Connectivity-guided targeting, programming, and EC evidence integration

Network-guided neuromodulation reframes targeting as the attempt to steer pathological dynamics through patient-specific connectivity structure. [136;137] In practice, this requires combining complementary connectivity classes: structural connectivity provides feasibility and safety constraints on which pathways are anatomically plausible; electrophysiological FC characterizes statistical coupling and its variation across vigilance and peri-ictal states; and EC, whether estimated from controlled perturbations or spontaneous propagation patterns, can add directional information that undirected association alone cannot provide. [38;118;138] Within the directionality boundary defined above, connectivity-guided target selection is most defensible when passive directional estimates are integrated with structural and perturbational evidence in a patient-specific decision framework.

Focal RNS programming illustrates this patient-specific principle in a clinical setting. For pathway engagement, patient-specific volume of tissue activated (VTA)-based tractography in mesial temporal RNS remained associated with seizure reduction after cross-validation, whereas normative healthy-subject connectivity did not. [139] For contact-level network evidence, larger early latency in-degree CCEP responses near eventual RNS contacts were associated with greater seizure reduction after RNS therapy. [140] This implication extends beyond RNS. Group-average or normative networks remain useful for hypothesis generation and syndrome-level context, but

programming-facing conclusions should be checked against within-patient evidence whenever possible. [140–143]

Within this patient-specific framework, programming decisions can draw on two complementary sources of directional connectivity evidence. Perturbational connectivity, including CCEP-derived connectivity, provides protocol-conditioned directional evidence from controlled stimulation rather than passive covariance, but its interpretation remains dependent on stimulation protocol, latency window, response definition, and sampling coverage. [38] TMS paired with iEEG provides a complementary perturbational mapping route, allowing non-invasive stimulation to be paired with intracranial recordings to characterize local and distributed evoked responses. At present, this evidence mainly supports mechanistic probing, connectivity mapping, and target-engagement validation, but it does not yet establish an epilepsy-specific role in stimulation programming or outcome-validated target selection. [144] Propagation-derived EC can add information about recruitment structure during interictal spike propagation; in source-space iEEG, information flow from spike onset to spread regions and resection of the spike-onset zone were associated with surgical outcome.[138] This route remains conditional on spike detection, source reconstruction, spatial sampling, and directionality-estimator assumptions. Together, these limitations make convergence across perturbational, propagation-derived, and structural evidence more informative for programming decisions than any single network label.

Across these evidence routes, the main deployment limitation is temporal: estimating a full connectivity matrix is generally poorly suited to millisecond-scale closed-loop stimulation unless causal relevance, low-latency feasibility, and chronic robustness are explicitly demonstrated. For programming-facing use, connectivity evidence is strongest when integrated with anatomical constraints, perturbational or propagation-based evidence, longitudinal stability, and validation against clinically meaningful outcomes.

7 Brain Modeling 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. [145;146] Evaluation should instead focus on issues that determine real use: what is 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. For brain models, this assessment also requires making explicit the plausible deployment scenario, real-time or near-real-time feasibility, patient-specific fitting burden, and integration burden within clinical or device workflows.[145;146]

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 10 organizes these model families into the same three deployment roles used in the text: monitoring and control-design representations, planning-oriented mechanistic and virtual models, and mapping or cross-modal transformation tools.

7.1 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 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 hidden Markov model comprising interictal, preictal, and ictal regimes, and later implementations combined hidden Markov models with time-frequency features for EEG-based detection and prediction.[147;148] 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 used state-space model identification to

represent EEG epochs with compact dynamical parameters for automated seizure detection. [149;150]

For monitoring and detection, the practical appeal of latent-state and state-space representations is that they reduce EEG observations or detector outputs to a small number of discrete states, continuous state values, or compact dynamical parameters that can be mapped onto Viterbi decoding, thresholding, classifier decisions, or alarm logic. [147–150] However, practical use depends on whether these state definitions remain interpretable and stable over time. The main risks are non-stationarity and concept drift across recording conditions, the need for retraining or recalibration as patient-specific state dynamics change, and limited robustness outside the monitoring conditions in which the models were calibrated. [78]

A related use of bounded representations is not only to monitor state transitions, but also to simplify the dynamics sufficiently for controller design and stimulation-policy testing. 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.[151;152] Data-driven reduced dynamics derived from intracranial recordings have similarly been proposed to support constrained stimulation design. [153]

For control design, these models are best viewed as abstractions. Their value lies in exposing simplified state variables, feedback laws, or locally fitted dynamical descriptions that can be analyzed for stability and tested in simulation without requiring a full patient-specific mechanistic reconstruction. The limiting issue is transfer; model mismatch, unobserved state changes, sparse iEEG sampling, and sensing- or stimulation-induced changes can all invalidate the fitted operating regime. Therefore, evidence that a surrogate supports controller design or proof-of-principle stimulation should be distinguished from evidence for autonomous chronic clinical control.[151–153]

7.2 Planning-oriented models rather than routine control objects

Greater biological detail does not imply greater clinical readiness. In this section, detailed models are evaluated primarily as planning tools unless they produce repeatedly observable, identifiable, and updatable variables for routine closed-loop use. [145;146]

7.2.1 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. [154] The Epileptor family has become a canonical example, providing a flexible seizure-generating formalism across multiple dynamical regimes. [155;156] 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.[157] For intervention, these models are appealing because they introduce explicit latent quantities, such as excitability or slow adaptation variables that govern transitions between normal and ictal states, suggesting 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. [158]

However, these latent quantities are not directly observed and are not uniquely constrained by routine clinical recordings. Better inference does not remove dependence on priors, model class, source-to-sensor mapping, and the completeness of SEEG sampling. The key translational issue is identifiability rather than mechanistic detail alone; fitted parameters or inferred epileptogenic-zone network (EZN) hypotheses should be reported with sensitivity analyses, posterior uncertainty, and explicit assumptions about the observation model. Computational cost remains relevant, because patient-specific model inversion commonly relies on repeated simulation, optimization, or posterior sampling, but the larger concern is whether the inferred

latent variables are sufficiently constrained to guide a patient-specific decision.[145;146;158;159]

7.2.2 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 the removal of nodes, regions, or connections and thereby estimate their modeled effect on network ictogenicity, synchronizability, or seizure propagation.[160;161] This makes them attractive for presurgical scenario testing and for prioritizing model-identified regions or connections that warrant closer clinical evaluation. Their validity depends on the network representation used to construct the model, whether structural or functional, the spatial coverage of the recorded or reconstructed network, and the dynamical assumptions used to define seizure propagation, synchronizability, or node ictogenicity. The implementation burden also differs across variants. Structural-network approaches require imaging-derived connectomes aligned to patient-specific anatomy; intracranial functional-network approaches depend on electrode localization, preprocessing, and annotated iEEG or ECoG recordings; and full virtual-brain workflows add measurement mapping, epileptogenic-zone priors, and patient-specific parameter inference. [145;146;160–162]

7.2.3 Personalized virtual-brain frameworks

Personalized virtual-brain frameworks extend network-level ictogenicity and virtual-resection logic by combining patient-specific anatomy, connectivity, dynamics, and measurement mappings into a single in-silico environment. [159;163] Existing syntheses emphasize their potential for diagnosis, surgery planning, and stimulation strategy design. [145;164] 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.[146;165] A recent stimulation-specific extension of this framework used high-resolution virtual brain twins to model SEEG and temporal-interference stimulation, estimate epileptogenic-zone networks from stimulation-induced seizures, and integrate SEEG and scEEG through source-to-sensor

mappings and Bayesian inversion. [166] A recent cohort-level study illustrates this benchmarking role: Dollomaja et al. generated thirty virtualized drug-resistant epilepsy patients from multimodal patient-specific imaging and SEEG recordings, with individualized epileptogenic-zone hypotheses and model parameters serving as synthetic ground truth for evaluating epileptogenic-zone estimation and source-localization methods rather than as clinical outcome validation. [159]

Nevertheless, even when fitted to empirical clinical recordings, these frameworks face uncertainty that accumulates across the modeling pipeline. Across imaging, connectivity estimation, source-to-sensor mapping, and parameter inversion, each step introduces its own uncertainty, and these uncertainties can propagate into the final simulated seizure spread, stimulation response, or epileptogenic-zone hypothesis. Different parameter configurations or epileptogenic-zone hypotheses may also reproduce similar SEEG patterns, requiring explicit treatment of posterior uncertainty and degeneracy. The translational value of personalized virtual-brain frameworks depends on whether EZN hypotheses, virtual-surgery scenarios, or stimulation strategies change patient-specific planning beyond existing clinical evidence and whether that added value is confirmed in prospective patient-specific validation.[145;146;158–160;166]

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.[146;167;168] Their value should therefore be judged not only by anatomical plausibility, but also by whether they improve decision-relevant inference under realistic uncertainty.[145;146]

Electrical source imaging (ESI) can add prospective diagnostic value in presurgical evaluation, including in magnetic resonance imaging (MRI)-negative epilepsy. [167;169;170] 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. For clinical translation, the question is whether the reconstruction improves

localization or planning decisions beyond sensor-space analysis and does so reproducibly. Its value lies in supporting localization and mapping under explicit uncertainty.

This validation problem becomes more demanding when the transformation crosses modalities or recording scales, because empirical ground truth is harder to establish and domain shift becomes a central deployment risk. 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.[14;168] Related approaches include distributed spatial filtering for intracranial arrays.[171] These directions are conceptually attractive because they aim to narrow the gap between what is sensed and what is required for intervention planning. However, cross-modal transformations 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 empirical route for testing whether sensor-level signatures correspond reproducibly to intracranial events in real recordings. [13]

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. Closed-loop systems are not simply detectors followed by stimulators. They are 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, then examine epilepsy-specific adaptation from clinician-governed programming to computational policy-learning frameworks, before

turning to the safety envelopes and practical clinical boundaries that define whether a policy can be regarded as clinically governable (Table 11).

8.1 Control variables and action bounds

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 program. 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. 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. [6;172] Clinical experience across epilepsy stimulation targets and patient groups similarly reinforces that long term outcomes depend on candidate selection and longitudinal calibration rather than stimulation delivery alone. [173–175]

RNS operationalized the canonical sensing-detection-stimulation loop for focal epilepsy as a reactive policy in which detected abnormal activity triggers tightly bounded stimulation bursts.[22] Long term follow-up shows that benefit can be durable over years, while also indicating that clinical optimization depends on longitudinal calibration and iterative programming rather than a fixed stimulation protocol. [23;176;177] Heterogeneity across patients and substrates, including mesial temporal lobe epilepsy, further reinforces that calibration and conservative operating points are part of deployment, not post-hoc programming details.[23;177;178] Recent SEEG-guided neocortical RNS data further suggest that implantation geometry is clinically relevant but insufficient to fully explain chronic response: proximity between the SEEG-defined focus and RNS depth leads predicted one-year response only among strong responders and did not remain significant at longer term follow-up, possibly reflecting contributions from network-level

adaptation over time. [179] Thalamic RNS is also emerging beyond conventional focal indications: retrospective pediatric and young-adult experience with centromedian or anterior thalamic leads suggests possible benefit in multifocal or regional-onset epilepsy, although the evidence remains small and caregiver-reported. [180] The NAUTILUS trial provides randomized sham-controlled evidence for responsive centromedian thalamic stimulation in drug-resistant idiopathic generalized epilepsy. Its prespecified endpoint based on time to the second generalized tonic-clonic seizure was not significant, but longer term and post-hoc analyses suggest clinically meaningful benefit with an acceptable safety profile.[181]

Deep brain stimulation of the anterior nucleus of the thalamus (ANT-DBS) provides a complementary reference because seizure reduction can arise through modulation of distributed networks rather than focal event suppression alone. [182] Long term registry and multicenter follow-up data further support ANT-DBS as a chronic clinical neuromodulation strategy, while indicating that response is shaped by epilepsy phenotype, prior surgical history, seizure severity, and programmable stimulation context rather than by target selection alone.[183;184] This widens the action space from immediate triggering toward scheduling, parameter selection, and clinician-guided mode changes. This action-space perspective is not specific to thalamic DBS. Comparative reviews across VNS, DBS, and RNS suggest that differences in targets and stimulation paradigms mainly change the available routes for personalization and adaptation, rather than defining wholly separate control principles. [185] Contemporary VNS registry evidence in Lennox-Gastaut syndrome shows sustained seizure reductions over 24 months with a low rate of treatment related adverse events, while also indicating that real-world benefit should be interpreted alongside dose escalation during follow-up and differences in response across seizure types rather than as the effect of a single fixed device setting.[186] However, in clinical devices, personalization is implemented by adjusting a limited set of programmable parameters rather than by allowing unrestricted control actions.[11;187] Representative stimulation parameters for RNS, VNS, and ANT-DBS are summarized in Table 12.

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.[11;17;188] Recent exploratory work on stimulating ictal fast-ripple hubs makes this trade-off concrete: fast-ripple-hub stimulation was associated with reduced seizure occurrence over a short inpatient post-stimulation window, and the same study reported a lower probability of stimulation-evoked seizures than seizure-onset-zone stimulation under matched stimulation intensities. However, these findings remain proof-of-concept, and translating such biomarkers into chronic closed-loop control would require evidence that the biomarker and derived network target can be calibrated, monitored, and remain stable over time. [189]

8.2 Governed adaptation in epilepsy

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.

8.2.1 Clinician-governed parameter adaptation

In current clinical practice, the most established form of adaptation in epilepsy is not autonomous online learning, but clinician-governed adjustment of bounded device parameters. In RNS, clinicians iteratively refine detection settings, sensing or stimulation pathways, and core stimulation parameters using stored ECoG, device detections, patient-reported seizures, clinical response, and tolerability. DBS programming is similarly individualized through repeated adjustment of stimulation settings, contact configuration, cycling schedule, and side-effect mitigation. This form of adaptation is

clinically important because epilepsy neuromodulation usually lacks immediate observable feedback after a programming change; clinicians therefore optimize a large parameter space using delayed and noisy seizure outcomes.[22;23;176;177;187] Recent longitudinal ANT-DBS sensing data illustrate how implanted recordings may support clinician-governed calibration: in a four-year observational cohort, higher-to-lower frequency power ratios in ANT local field potentials helped distinguish responders from non-responders and were proposed as candidate indicators for physiology-guided programming optimization, while circadian and multidien rhythms were noted as possible inputs for future state-dependent strategies; these directions remain to be tested prospectively.[190] In generalized epilepsy, preliminary intraoperative RNS data from the centromedian nucleus of the thalamus (CM) region, acquired under anesthesia from a three-patient cohort, suggest that stronger stimulation-evoked delta responses at low-activity CM-adjacent sites, likely involving the internal medullary lamina, may warrant further study as a target-refinement signal.[191]

8.2.2 Retrospective state-dependent evidence

Exploratory retrospective analyses of chronic RNS data have begun to suggest candidate state-dependent timing and parameter-selection principles. Chiang et al. showed that the effects of RNS parameters on seizure-risk transitions depend on the current risk state, while Anderson et al. reported that a higher proportion of stimulation during low-risk epochs early in therapy was associated with greater seizure reduction. [192;193] Together with evidence that seizures are modulated by circadian and multidien rhythms, these findings broaden the closed-loop question from a narrow detect-and-abort paradigm toward longer-timescale risk-state modulation.[194] However, because these studies are retrospective and rely on device-derived summaries or inferred risk states, they should be treated as hypothesis-generating rather than as validated adaptive stimulation policies.

8.2.3 Computational and preclinical policy learning

A separate body of computational and preclinical work has made the policy-optimization problem more explicit. Batch reinforcement-learning studies have used in

vitro epileptiform recordings from animal brain-slice preparations to learn adaptive stimulation policies from labeled state-action trajectories. Koopman-operator and model-predictive-control frameworks have combined real or synthetic EEG prediction with seizure-suppression experiments in neural-mass simulation models, whereas recent deep reinforcement-learning controllers have been evaluated in computational cortico-thalamic seizure models. Because these studies remain grounded in preclinical preparations, recorded EEG prediction tasks, or computational seizure models, they are best interpreted as algorithmic feasibility evidence rather than as clinically validated adaptive neuromodulation policies. [195–198]

Taken together, these evidence classes suggest that adaptive DBS in Parkinson's disease is best treated as an engineering comparator rather than as direct epilepsy evidence, illustrating how low-dimensional biomarkers can be mapped to bounded stimulation actions under explicit safety, latency, energy, artifact, and calibration constraints. [17;188;199] 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. [11;200;201] 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.[200;201] The common implication is a governance boundary. Learning may support a clinician-defined control surface, but autonomous policy revision still lacks the prospective stability and safety evidence required for epilepsy deployment.

Accordingly, 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, monitored, and clinically governed. 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. For closed-loop and adaptive stimulation, the evaluation framework (Table 4) treats retrospective benefit or offline risk-state discrimination as insufficient without prospective evidence of stability, recalibration burden, clinician oversight, and safety related failure modes.

8.3 Safety envelopes and governance boundaries

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. Operationally, partial observability can be bounded by reporting the spatial coverage of recording contacts relative to presumed seizure onset and propagation sites, the fraction of clinically relevant events captured by the sensing montage, and the longitudinal stability and uncertainty of the biomarker-to-state mapping. [11;17;179] This framing helps distinguish differences in implementation and clinical objective from differences in underlying control principle. The immediate translational implication is not that epilepsy stimulation should pursue increasingly complex controllers, but that closed-loop policies must make their safety limits, stimulation bounds, artifact-handling assumptions, calibration procedures, and clinician-supervision points explicit.[17;188]

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.[202] Conceptual analyses likewise emphasize that stimulation effects are network and state-dependent, cautioning against simplistic mechanistic narratives and motivating context-aware policies.[203] 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. [204] 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.

A closed-loop policy cannot be considered clinically governable unless its action space is bounded, its major signal and stimulation failure modes are explicit, conservative behavior is predefined when the feedback signal is unavailable, excessively contaminated by artifact, or too uncertain to support reliable triggering, and clinician supervision is preserved. End-to-end latency, computational and power budget, stimulation-artifact handling, charge-balance, charge-density and stimulation-duration or duty-cycle constraints, drift monitoring, and the clinician-facing control surface should be treated as part of the policy itself rather than as reporting formalities.[11;17;188] For adaptive policies, the same safety envelope should also specify which detection or stimulation parameters are eligible for update, the allowed update cadence or triggering conditions, and the conditions under which adaptation is suspended or returned to a predefined clinician-approved setting. [11;17;188] In practice, these conditions define the minimum boundary between a governable clinical system and a research prototype.

9 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 for 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.[205] 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. [8]

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. Uncertainty estimation, abstention mechanisms, and patient-level monitoring of calibration drift should be incorporated into longitudinal 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.[8] Box 2 summarizes these reporting requirements as a clinical deployment evaluation checklist.

Box 2. Clinical deployment evaluation checklist

Synthesized from Table 4, this checklist groups reporting metrics by validation level.

Discrimination: area under the curve (AUC), sensitivity, precision, event-level detection rate, and false positives per unit time.

Temporal: sensitivity at a fixed false-alarm budget, calibration drift, chronological train-test split, and leakage control.

Prospective: processing latency, time-in-warning, abstention rate, uncertainty calibration, and clinician review burden.

Chronic: long term degradation curve, recalibration frequency, operating-point stability, override rate, safety failure modes, and rollback events.

9.2 Control safety and adaptive governance

Another 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. 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. 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

The last 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.

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 Limitations of the review design

This review has several limitations that should be acknowledged. Although a structured literature search was conducted with explicit search strings, inclusion criteria and database coverage, the synthesis remains a scoping rather than exhaustive review. The scope is further bounded by a focus on algorithmic and signal-processing approaches. Hardware-level implementations, including custom integrated circuits, neuromorphic chips, and embedded systems designed for on-device EEG processing or stimulation delivery, are not covered, even where such implementations are clinically relevant. Priority given to clinically and translationally relevant evidence also means that some task-specific algorithmic studies and purely offline methodological comparisons are not discussed in detail.

Furthermore, the broader literature on which this review draws reflects the current state of the field. Although long term ambulatory and wearable validation studies are emerging, studies that simultaneously combine prospective design, multicenter recruitment, long term monitoring, and operation outside controlled research settings remain uncommon. Most available evidence derives from retrospective datasets, short recordings, single-center cohorts, or device-specific implementations, which constrains the strength of conclusions that any review of this topic can currently offer regarding long term reliability, calibration stability, adaptive control, and clinical integration.

11 Conclusion

Read across the sensing-to-control pipeline, a common pattern emerges: deployment failures often arise less from a single weak algorithm than from uncertainty that is passed forward across stages. Recording modality, spatial and temporal sampling, preprocessing choices, artifact handling, event definitions, connectivity estimates, and model assumptions each constrain what later stages can validly infer or act upon. The central deployment task is to make these handoffs explicit, so that downstream programming or closed-loop decisions are not built on unexamined assumptions.

Building on the validation boundaries outlined above, future studies should prioritize prospective, longitudinal, and workflow-aware evaluation under realistic clinical conditions. The critical test is whether operating points and uncertainty estimates remain stable, and whether the clinical decisions they support remain reliable, when methods are evaluated chronologically, in patient-specific contexts, and within the clinical workflows in which they are intended to be used.

For adaptive neuromodulation, the most credible near term path is governed adaptation rather than unrestricted autonomy. This requires stable state summaries, bounded action spaces, clinician-supervision checkpoints, and auditable criteria for escalation to clinician review, rollback to a prior approved setting, or suspension of adaptation when signals drift or evidence becomes unreliable. Ultimately, clinically durable systems will be judged by whether they remain safe, interpretable, and useful over the years-long timescales that matter to patients and clinicians.

List of Abbreviations

AI Artificial intelligence

ANT Anterior nucleus of the thalamus

ANT-DBS Anterior nucleus of the thalamus deep brain stimulation

AUC Area under the ROC curve

BDSP Brain Data Science Platform

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)

DBS Deep brain stimulation

DL Deep learning

DR Delayed response

DTI Diffusion tensor imaging

DWI Diffusion-weighted imaging

EC Effective connectivity

ECoG Electrocorticography

EDF European Data Format

EEG Electroencephalography

EMG Electromyography

EMU Epilepsy monitoring unit

ER Early response

ESI Electrical source imaging

EZN Epileptogenic zone network

FC Functional connectivity

fMRI Functional magnetic resonance imaging

FP/min False positives per minute

hd-EEG High-density EEG

hd-scEEG High-density scalp EEG

HEEDB Harvard Electroencephalography Database

HFO High-frequency oscillation

HMM Hidden Markov model

ICA Independent component analysis

ICU Intensive care unit

IED Interictal epileptiform discharge

iEEG Intracranial EEG

iEEG-BIDS Intracranial EEG extension of BIDS

IEEE Institute of Electrical and Electronics Engineers

LFP Local field potential

MEG Magnetoencephalography

ML Machine learning

MRI Magnetic resonance imaging

MTL Medial temporal lobe

N1 Negative component 1

OSF Open Science Framework

PET Positron emission tomography

PPV Positive predictive value

PRISMA-ScR Preferred Reporting Items for Systematic Reviews and Meta-Analyses Extension for Scoping Reviews

RNS Responsive neurostimulation

ROC Receiver operating characteristic

SC Structural connectivity

scEEG Scalp EEG

SEEG Stereo-electroencephalography

SIR Susceptible-infected-recovered model

SNR Signal-to-noise ratio

SOZ Seizure onset zone

SPES Single-pulse electrical stimulation

spkHFO High-frequency oscillation co-occurring with spike

tACS Transcranial alternating current stimulation

tES Transcranial electrical stimulation

TMS Transcranial magnetic stimulation

TUH Temple University Hospital

TUSZ Temple University Hospital EEG Seizure Corpus

VEP Virtual Epileptic Patient

VNS Vagus nerve stimulation

VTA Volume of tissue activated

wPLI Weighted phase lag index

Author contributions

Zhuoran Xu, Sepehr Shirani, Antonio Valentín Huete, and Ioannis Stavropoulos conceived the study. Zhuoran Xu, Sepehr Shirani, Saeid Sanei, Antonio Valentín Huete,

and Ioannis Stavropoulos designed the study. Zhuoran Xu, Sepehr Shirani, Saeid Sanei, Gonzalo Alarcón, Antonio Valentín Huete, and Ioannis Stavropoulos defined and interpreted the intellectual content. Zhuoran Xu, Sepehr Shirani, and Jiaxin Lei performed the literature search. Zhuoran Xu drafted the manuscript. Zhuoran Xu, Sepehr Shirani, Saeid Sanei, Gonzalo Alarcón, Antonio Valentín Huete, and Ioannis Stavropoulos critically revised the manuscript. Antonio Valentín Huete and Ioannis Stavropoulos served as guarantors of the work. All authors reviewed and approved the final manuscript.

Conflict of Interest

The authors declare no conflicts of interest.

Declaration of AI and AI-assisted technologies in the writing process

During the preparation and revision of this work, the authors used ChatGPT (OpenAI) and Claude (Anthropic) to assist with language refinement and consistency checks in selected parts of the manuscript. The authors retained responsibility for the scientific content, interpretation, and conclusions. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Data Availability Statement

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

References

1. Feigin VL, Vos T, Nair BS, et al. Global, regional, and national burden of epilepsy, 1990–2021: a systematic analysis for the Global Burden of Disease Study 2021. Lancet Public Health. 2025;10(3):e203–e227.

2. Perucca E, Perucca P, White HS, Wirrell EC. Drug resistance in epilepsy. Lancet Neurol. 2023;22(8):723–734.
3. 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.
4. Lucas A, Revell A, Davis KA. Artificial intelligence in epilepsy—applications and pathways to the clinic. Nat Rev Neurol. 2024;20(6):319–336.
5. 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.
6. Daraie AH, Damiani A, Khoshkhou M, et al. Artificial intelligence for adaptive neuromodulation in drug-resistant epilepsy. Epilepsia. 2026:1–16.
7. Ewen JB, Babiloni C, Collins GS, et al. The GREENBEAN checklist for reporting studies evaluating the effectiveness of EEG-based biomarkers. Clin Neurophysiol. 2025;176:2110777.
8. 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.
9. Segal G, Keidar N, Lotan RM, Romano Y, Herskovitz M, Yaniv Y. Utilizing risk-controlling prediction calibration to reduce false alarm rates in epileptic seizure prediction. Front Neurosci. 2023;17:1184990.
10. 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.
11. Kassiri H, Muneeb A, Salahi R, Dabbaghian A. Closed-loop implantable neurostimulators for individualized treatment of intractable epilepsy: A review of recent developments, ongoing challenges, and future opportunities. IEEE Trans Biomed Circuits Syst. 2024.
12. Li H, Wang Y, Tanabe S, et al. Mapping epileptic directional brain networks using intracranial EEG data. Biostatistics. 2021;22(3):613–628.
13. 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.
14. 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. 2025;25(2):494.
15. 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.

16. 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.
17. Stanslaski S, Summers RL, 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.
18. Goldberger AL, Amaral LA, Glass L, et al. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. Circulation. 2000;101(23):e215–e220.
19. 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.
20. 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.
21. Al-Sadek T, Wadhwa A, Wadhwa M, Warren AE, Rolston JD. Methodologies to detect cortico-cortical evoked potentials: a systematic review. Front Hum Neurosci. 2025;19:1636115.
22. Morrell MJ. Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. Neurology. 2011;77(13):1295–1304.
23. 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.
24. 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.
25. 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.
26. Parmigiani S, Mikulan E, Russo S, et al. Simultaneous stereo-EEG and high-density scalp EEG recordings to study the effects of intracerebral stimulation parameters. Brain Stimul. 2022;15(3):664–675.
27. Van Blooijs D, Blok S, Huiskamp GJ, Van Eijsden P, Meijer HG, Leijten F. The effect of propofol on effective brain networks. Clin Neurophysiol. 2024;161:222–230.
28. Jobst BC, Bartolomei F, Diehl B, et al. Intracranial EEG in the 21st century. Epilepsy Curr. 2020;20(4):180–188.
29. 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.
30. Uriguen JA, Garcia-Zapirain B. EEG artifact removal—state-of-the-art and guidelines. J Neural Eng. 2015;12(3):031001.

31. 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.
32. 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.
33. 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.
34. Rubboli G, Bø MH, Alfstad K, et al. Clinical utility of ultra long-term subcutaneous electroencephalographic monitoring in drug-resistant epilepsies: a “real world” pilot study. Epilepsia. 2024;65(11):3265–3278.
35. Munteanu T, Daly B, O'Donnell C, et al. Unilateral ultra long-term subcutaneous EEG monitoring in drug-refractory idiopathic generalized epilepsy. Epilepsia. 2026;67(1):396–407.
36. Nasseri M, Stirling RE, Viana PF, et al. Forecasting epileptic seizures with wearable devices: A hybrid short- and long-horizon pseudo-prospective approach. Epilepsia. 2025.
37. Khambhati AN, Chang EF, Baud MO, Rao VR. Hippocampal network activity forecasts epileptic seizures. Nat Med. 2024;30(10):2787–2790.
38. Matsumoto R, Kunieda T, Nair D. Single pulse electrical stimulation to probe functional and pathological connectivity in epilepsy. Seizure. 2017;44:27–36.
39. Obeid I, Picone J. The temple university hospital EEG data corpus. Front Neurosci. 2016;10:196.
40. Shah V, Von Weltin E, Lopez S, et al. The temple university hospital seizure detection corpus. Front Neuroinform. 2018;12:83.
41. 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.
42. Klatt J, Feldwisch-Drentrup H, Ihle M, et al. The EPILEPSIAE database: An extensive electroencephalography database of epilepsy patients. 2012.
43. Karoly PJ, Ung H, Grayden DB, et al. The circadian profile of epilepsy improves seizure forecasting. Brain. 2017;140(8):2169–2182.
44. 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.
45. Falach R, Geva-Sagiv M, Eliashiv D, et al. Annotated interictal discharges in intracranial EEG sleep data and related machine learning detection scheme. Sci Data. 2024;11(1):1354.
46. Duan C, Zhang Y, Kanai S, et al. Omni-iEEG: A Large-Scale, Comprehensive iEEG Dataset and Benchmark for Epilepsy Research. in The Fourteenth International Conference on Learning Representations 2026.

47. Kini LG, Davis KA, Wagenaar JB. Data integration: Combined imaging and electrophysiology data in the cloud. Neuroimage. 2016;124:1175–1181.
48. 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.
49. Sharma R, Meena HK. Emerging trends in EEG signal processing: A systematic review. SN Comput Sci. 2024;5(4):415.
50. Chaddad A, Wu Y, Kateb R, Bouridane A. Electroencephalography signal processing: A comprehensive review and analysis of methods and techniques. Sensors. 2023;23(14):6434.
51. Widmann A, Schröger E, Maess B. Digital filter design for electrophysiological data–a practical approach. J Neurosci Methods. 2015;250:34–46.
52. Bénar CG, Chauvière L, Bartolomei F, Wendling F. Pitfalls of high-pass filtering for detecting epileptic oscillations: a technical note on “false” ripples. Clin Neurophysiol. 2010;121(3):301–310.
53. 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.
54. 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):1–6 IEEE 2022.
55. Lopes F, Leal A, Medeiros J, et al. Automatic electroencephalogram artifact removal using deep convolutional neural networks. IEEE Access. 2021;9:149955–149970.
56. Zhou A, Johnson BC, Muller R. Toward true closed-loop neuromodulation: artifact-free recording during stimulation. Curr Opin Neurobiol. 2018;50:119–127.
57. Mouthaan BE, Klooster MA, Keizer D, et al. Single Pulse Electrical Stimulation to identify epileptogenic cortex: Clinical information obtained from early evoked responses. Clin Neurophysiol. 2016;127(2):1088–1098.
58. 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.
59. Kohli S, Casson AJ. Removal of gross artifacts of transcranial alternating current stimulation in simultaneous EEG monitoring. Sensors. 2019;19(1):190.
60. 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.
61. Varone G, Hussain Z, Sheikh Z, et al. Real-time artifacts reduction during TMS-EEG coregistration: a comprehensive review on technologies and procedures. Sensors. 2021;21(2):637.

62. 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.
63. 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.
64. Retana M, 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.
65. Chen W, Liu X, Wan P, Chen Z, Chen Y. Anti-artifacts techniques for neural recording front-ends in closed-loop brain-machine interface ICs. Front Neurosci. 2024;18:1393206.
66. Liu TC, Chen PL, Chen YC, et al. Efficient artifact removal for adaptive deep brain stimulation and a temporal event localization analysis. J Neurosci Methods. 2026:110708.
67. Prime D, Woolfe M, O'Keefe S, Rowlands D, Dionisio S. Quantifying volume conducted potential using stimulation artefact in cortico-cortical evoked potentials. J Neurosci Methods. 2020;337:108639.
68. Levinson L, Sun S, Paschall C, et al. Data processing techniques impact quantification of cortico-cortical evoked potentials. J Neurosci Methods. 2024;408:110130.
69. 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.
70. Sudhakar JG, Elango K. A Comprehensive Review of EEG-Based Seizure Detection Techniques. IEEE Access. 2025.
71. Deng Z, Li C, Zhao G, Chen X. Incremental Learning for Patient-Specific EEG-Based Seizure Detection. IEEE Trans Neural Syst Rehabil Eng. 2025.
72. 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.
73. 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.
74. Andrzejak RG, Zaveri HP, Schulze-Bonhage A, et al. Seizure forecasting: Where do we stand?. Epilepsia. 2023;64:S62–S71.
75. 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.
76. 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.
77. Lopes F, Leal A, Pinto MF, et al. Removing artefacts and periodically retraining improve performance of neural network-based seizure prediction models. Sci Rep. 2023;13(1):5918.

78. Pontes ED, Pinto M, Lopes F, Teixeira C. Concept-drifts adaptation for machine learning EEG epilepsy seizure prediction. Sci Rep. 2024;14(1):8204.
79. 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.
80. Viana PF, Pal Attia T, Nasseri M, et al. Seizure forecasting using minimally invasive, ultra-long-term subcutaneous electroencephalography: individualized intrapatient models. Epilepsia. 2023;64:S124–S133.
81. Yang H, Mueller J, Eberlein M, et al. Seizure forecasting with ultra long-term EEG signals. Clin Neurophysiol. 2024;167:211–220.
82. Sun Y, Guan M, Chen X, et al. Deep learning-based classification and segmentation of interictal epileptiform discharges using multichannel electroencephalography. Epilepsia. 2025.
83. 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.
84. 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.
85. 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.
86. 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.
87. 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.
88. Park S, Clark JS, Viana PF, et al. Seizure detection using ultra-long-term subcutaneous electroencephalography: A deep learning CNN-BiLSTM approach. Epilepsia. 2026;67(1):437–447.
89. Karoly PJ, Cook MJ, Maturana M, et al. Forecasting cycles of seizure likelihood. Epilepsia. 2020;61(4):776–786.
90. Yang H, Müller J, Kalousios S, et al. Seizure forecasting with epilepsy cycles: On the causality of forecasting pipelines. Epilepsia. 2026.
91. 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.
92. Stirling RE, Cook MJ, Grayden DB, Karoly PJ. Seizure forecasting and cyclic control of seizures. Epilepsia. 2021;62:S2–S14.
93. Xiong W, Stirling RE, Payne DE, et al. Forecasting seizure likelihood from cycles of self-reported events and heart rate: a prospective pilot study. EBioMedicine. 2023;93.

94. 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.
95. Ding TY, Gagliano L, Jahani A, Toffa DH, Nguyen DK, Bou Assi E. Epileptic seizure forecasting with wearable-based nocturnal sleep features. Epilepsia Open. 2024;9(5):1793–1805.
96. 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.
97. 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.
98. Jiahao H, Rahman MMU, Laleg-Kirati TM. Calibrated Uncertainty Estimation based on Bayesian Models for EEG Epileptic Seizure Detection. IEEE Sens J. 2026.
99. 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.
100. 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.
101. 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.
102. 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.
103. 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.
104. 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.
105. 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.
106. Qiang Z, Norris J, Cooray G, et al. Cortico-cortical evoked potentials: Analytical techniques and emerging paradigms for epileptogenic zone localization. Epilepsia. 2025.
107. 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.
108. 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.
109. 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.

110. 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.
111. 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.
112. 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.
113. 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.
114. 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.
115. 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.
116. 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.
117. 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.
118. Friston KJ. Functional and effective connectivity: a review. Brain Connect. 2011;1(1):13–36.
119. Yaffe RB, Borger P, Megevand P, et al. Physiology of functional and effective networks in epilepsy. Clin Neurophysiol. 2015;126(2):227–236.
120. Parker CS, Clayden JD, Cardoso MJ, et al. Structural and effective connectivity in focal epilepsy. Neuroimage Clin. 2018;17:943–952.
121. Vinck M, Oostenveld R, Van Wingerden M, Battaglia F, Pennartz CM. 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.
122. Bastos AM, Schoffelen JM. A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. Front Syst Neurosci. 2016;9:175.
123. Seth AK, Barrett AB, Barnett L. Granger Causality Analysis in Neuroscience and Neuroimaging. J Neurosci. 2015;35(8):3293–3297.
124. Rubinov M, Sporns O. Complex network measures of brain connectivity: uses and interpretations. Neuroimage. 2010;52(3):1059–1069.
125. Van Wijk BC, Stam CJ, Daffertshofer A. Comparing brain networks of different size and connectivity density using graph theory. PLoS One. 2010;5(10):e13701.

126. Shi H, Pattnaik AR, Aguila C, et al. Utility of intracranial EEG networks depends on rereferencing and connectivity choice. Brain Commun. 2024;6(3):fcae165.
127. Van Mierlo P, Höller Y, Focke NK, Vulliemoz S. Network perspectives on epilepsy using EEG/MEG source connectivity. Front Neurol. 2019;10:721.
128. Chinedu-Eneh EO, Chiang S, Andrews JP, et al. Influences of electrode density on intracranial seizure localisation: a single-blinded randomised crossover study. EBioMedicine. 2025;113.
129. 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.
130. Royer J, Bernhardt BC, Larivière S, et al. Epilepsy and brain network hubs. Epilepsia. 2022;63(3):537–550.
131. Balatskaya A, Roehri N, Lagarde S, et al. The “Connectivity Epileptogenicity Index” (cEI), a method for mapping the different seizure onset patterns in StereoElectroEncephalography recorded seizures. Clin Neurophysiol. 2020;131(8):1947–1955.
132. 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.
133. Cousyn L, Messaoud RB, Lehongre K, et al. Daily resting-state intracranial EEG connectivity for seizure risk forecasts. Epilepsia. 2023;64(2):e23–e29.
134. 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.
135. 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.
136. Piper RJ, Richardson RM, Worrell G, et al. Towards network-guided neuromodulation for epilepsy. Brain. 2022;145(10):3347–3362.
137. 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.
138. 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.
139. 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.
140. 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.

141. 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.
142. 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.
143. Gregg NM, Ojeda Valencia G, Pridalova T, et al. Thalamic stimulation induced changes in network connectivity and excitability in epilepsy. Ann Neurol. 2025.
144. Tsang EW, Tatz JR, Hassan U, et al. Transcranial magnetic stimulation with intracranial recording in humans and primates: a review. Brain. 2026;149(5):1498–1513.
145. Jirsa V, Wang H, Triebkorn P, et al. Personalised virtual brain models in epilepsy. Lancet Neurol. 2023;22(5):443–454.
146. 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.
147. 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.
148. Abdullah MH, Abdullah JM, Abdullah MZ. Seizure detection by means of hidden Markov model and stationary wavelet transform of electroencephalograph signals. in Proceedings of 2012 IEEE-EMBS International Conference on Biomedical and Health Informatics:62–65 IEEE 2012.
149. 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.
150. Wang Z, Sperling MR, Wyeth D, Guez A. Automated seizure detection based on state-space model identification. Sensors. 2024;24(6):1902.
151. 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.
152. Chakravarthy N, Sabesan S, Tsakalis K, Iasemidis L. Controlling epileptic seizures in a neural mass model. J Comb Optim. 2009;17(1):98–116.
153. Ashourvan A, Pequito S, Khambhati AN, et al. Model-based design for seizure control by stimulation. J Neural Eng. 2020;17(2):026009.
154. Jirsa VK, Stacey WC, Quilichini PP, Ivanov AI, Bernard C. On the nature of seizure dynamics. Brain. 2014;137(8):2210–2230.
155. El Houssaini K, Ivanov AI, Bernard C, Jirsa VK. Seizures, refractory status epilepticus, and depolarization block as endogenous brain activities. Phys Rev E. 2015;91(1):010701.
156. 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).
157. 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.

158. 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.
159. 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.
160. Kini LG, Bernabei JM, Mikhail F, et al. Virtual resection predicts surgical outcome for drug-resistant epilepsy. Brain. 2019;142(12):3892–3905.
161. 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.
162. Junges L, Lopes MA, Terry JR, Goodfellow M. The role that choice of model plays in predictions for epilepsy surgery. Sci Rep. 2019;9(1):7351.
163. Wang HE, Triebkorn P, Breyton M, et al. Virtual brain twins: from basic neuroscience to clinical use. Natl Sci Rev. 2024;11(5):nwae079.
164. 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.
165. Jirsa VK, Proix T, Perdikis D, et al. The virtual epileptic patient: individualized whole-brain models of epilepsy spread. Neuroimage. 2017;145:377–388.
166. Wang HE, Dollomaja B, Triebkorn P, et al. Virtual brain twins for stimulation in epilepsy. Nat Comput Sci. 2025:1–15.
167. 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.
168. 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.
169. 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.
170. Sharma P, Scherg M, Pinborg LH, et al. Ictal and interictal electric source imaging in presurgical evaluation: a prospective study. Eur J Neurol. 2018;25(9):1154–1160.
171. 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):1117–1121 IEEE 2024.
172. Ohlerth AK, Valentin A, Vergani F, Ashkan K, Bastiaanse R. The verb and noun test for perioperative testing (VAN-POP): standardized language tests for navigated transcranial magnetic stimulation and direct electrical stimulation. Acta Neurochir. 2020;162(2):397–406.

173. 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.
174. 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.
175. Valentín A, Selway RP, Amarouche M, et al. Intracranial stimulation for children with epilepsy. Eur J Paediatr Neurol. 2017;21(1):223–231.
176. 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.
177. Rao VR. Personalizing Responsive Neurostimulation for Epilepsy. J Clin Neurophysiol. 2025;42(6):505–512.
178. 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.
179. Sadeghzadeh S, Purger DA, Bhanot P, et al. sEEG-guided responsive neurostimulation to treat neocortical epilepsy: A multicenter retrospective study of the efficacy and safety of depth electrode-mediated neuromodulation. Epilepsia Open. 2026;11(1):146–161.
180. Mohanty D, Houck KM, Trandafir C, et al. Responsive neurostimulation of thalamic nuclei for regional and multifocal drug-resistant epilepsy in children and young adults. J Neurosurg Pediatr. 2024;34(1):40–48.
181. Uysal U, Landazuri P, Burdette DE, et al. Responsive stimulation of the thalamus for idiopathic generalized epilepsy: Results of the randomized controlled NAUTILUS trial through 18 months. Epilepsia. 2026.
182. 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.
183. Kaufmann E, Peltola J, Colon AJ, et al. Long-term evaluation of anterior thalamic deep brain stimulation for epilepsy in the European MORE registry. Epilepsia. 2024;65(8):2438–2458.
184. Hu T, Xie H, Shan M, et al. Long-term efficacy of anterior nucleus of the thalamus deep brain stimulation in drug-resistant epilepsy: a 5-year multicenter study. BMC Med. 2025;23(1):1–14.
185. 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.
186. Lyons P, Wheless J, Verner R, et al. Vagus nerve stimulation in Lennox–Gastaut syndrome: Twenty-four-month data and experience from the CORE-VNS study. Epilepsia. 2025;66(5):1540–1549.
187. Wadi L, Pati S, Agashe S. Individualizing Programming of Responsive Neurostimulation and Deep Brain Stimulation Therapies in Epilepsy. J Clin Neurophysiol. 2025;42(6):513–520.

188. 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.
189. 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.
190. Aiello G, Elshahabi A, Stieglitz L, et al. Intracranial biomarkers for anterior thalamic deep brain stimulation in epilepsy: a long-term observational study. Brain. 2026:awag120.
191. Ho JC, Aung T, Damiani A, et al. Refining centromedian nucleus stimulation for generalized epilepsy with targeting and mechanistic insights from intraoperative electrophysiology. Nat Commun. 2025;16(1):5272.
192. Chiang S, Khambhati AN, Wang ET, Vannucci M, Chang EF, Rao VR. Evidence of state-dependence in the effectiveness of responsive neurostimulation for seizure modulation. Brain Stimul. 2021;14(2):366–375.
193. Anderson DN, Charlebois CM, Smith EH, et al. Closed-loop stimulation in periods with less epileptiform activity drives improved epilepsy outcomes. Brain. 2024;147(2):521–531.
194. Baud MO, Kleen JK, Mirro EA, et al. Multi-day rhythms modulate seizure risk in epilepsy. Nat Commun. 2018;9(1):88.
195. Pineau J, Guez A, Vincent R, Panuccio G, Avoli M. Treating epilepsy via adaptive neurostimulation: a reinforcement learning approach. Int J Neural Syst. 2009;19(04):227–240.
196. Guez A, Vincent RD, Avoli M, Pineau J. Adaptive Treatment of Epilepsy via Batch-mode Reinforcement Learning. in AAAI;8:1671–1678 2008.
197. Liang Z, Luo Z, Liu K, Qiu J, Liu Q. Online learning koopman operator for closed-loop electrical neurostimulation in epilepsy. IEEE J Biomed Health Inform. 2022;27(1):492–503.
198. Dan R, Zhang H, Bai J. Closed-Loop Control of Epilepsy Based on Reinforcement Learning. Int J Neural Syst. 2025;36(2):2550074–2550074.
199. Little S, Pogosyan A, Neal S, et al. Adaptive deep brain stimulation in advanced Parkinson disease. Ann Neurol. 2013;74(3):449–457.
200. Ortiz-Guerrero G, Gregg NM. Biomarkers for Epilepsy Deep Brain Stimulation. J Clin Neurophysiol. 2025;42(6):486–492.
201. Okoroafor F, Qiang Z, Rosenke S, et al. Neurophysiologic Biomarkers of Invasive Neuromodulation Therapy for Epilepsy. Neuromodulation. 2026.
202. Guidotti R, Basti A, Pieramico G, et al. When neuromodulation met control theory. J Neural Eng. 2025;22(1):011001.
203. Trevelyan AJ, Marks VS, Graham RT, Denison T, Jackson A, Smith EH. On brain stimulation in epilepsy. Brain. 2025;148(3):746–752.

204. Fröhlich F, McCormick DA. Endogenous electric fields may guide neocortical network activity. Neuron. 2010;67(1):129–143.
205. 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.

Tables 1-12 and Supplementary Table S1

Table 1. Operational four axis framework for comparing neuromodulation pipeline stages.

Pipeline stage	Temporal scale / latency	Observability / representation	Non-stationarity / drift	Deployment role
Acquisition and datasets	Report recording duration, sampling rate, follow-up length, annotation density, and whether data support offline, quasi-online, or prospective evaluation.	Report modality, sensor density, anatomical coverage, reference scheme, label source, and whether variables are sensor-level, source-level, event-level, or latent-state representations.	Report dataset shift, inter-patient heterogeneity, cross-session variability, medication or sleep-state information, and longitudinal changes in recording quality or event statistics.	Defines the upper bound on feasible inference, generalization, calibration, and clinical claims supported by the available data.
Preprocessing and artifact suppression	Report filter order, passband, phase properties, analysis-window length, causal versus non-causal implementation, and processing latency where online use is claimed.	Report how rereferencing, filtering, denoising, and artifact rejection alter the retained sensor-level signal and whether pathological activity may be attenuated or distorted.	Report sensitivity to reference choice, stimulation artifacts, impedance changes, movement, muscle activity, sleep state, and protocol or hardware variation.	Defines the effective observability boundary for all downstream sensing, modeling, and control-relevant inference.
Spontaneous events and evoked responses	Report event duration, detection latency, analysis-window length, warning horizon, stimulation-response window, and whether inference can be performed causally.	Report whether outputs are sensor-level events, source-localized events, risk states, stimulation-evoked response summaries where evoked responses are analyzed, or latent variables inferred from event structure.	Report threshold drift, calibration burden, cross-session stability, state dependence, medication effects, and sensitivity to stimulation timing or background activity.	Provides monitoring, warning, mapping, or control-trigger variables; requires event-level sensitivity together with false positives per unit time and temporal validation.
Connectivity analysis	Report window length, update interval, aggregation timescale, computational delay, and whether estimates can be updated online or only offline.	Report whether structural, functional, or effective connectivity is estimated; specify node definition, edge estimator, reference effects, leakage correction, and sampling coverage.	Report state dependence, windowing sensitivity, common-source or volume-conduction effects, missing-node bias, and stability across sessions or stimulation conditions.	Supports mapping, target ranking, stratification, or decision support; real-time claims require causal estimation, stable operating points, and interpretable uncertainty.
State estimation and brain models	Report inference interval, model-update frequency, planning horizon, computational cost, and whether recursive or offline estimation is used.	Report observed variables, latent-state definition, source model, parameter identifiability, uncertainty, and which physiological variables remain unobserved.	Report model mismatch, parameter drift, stimulation-response changes, recalibration needs, and robustness to chronic evolution of patient state or recording conditions.	Provides state summaries for monitoring, simulation, treatment planning, and policy support; stronger claims require uncertainty reporting and longitudinal validation.
Closed-loop control	Report sensing-to-stimulation latency, decision interval, stimulation-trigger delay, adaptation interval, and separation between fast triggering and slow clinician-governed tuning.	Report policy input variables, stimulation action space, safety constraints, uncertainty handling, and whether control depends on directly sensed events or inferred latent states.	Report policy drift, calibration decay, stimulation artifacts, stimulation-induced plasticity, false-trigger behavior, and failure modes under chronic use.	Implements constrained stimulation under clinician oversight; deployment claims require latency, safety envelope, override or rollback mechanisms, and prospective or chronic reliability evidence.

Note: Each axis is operationalized as a set of reportable engineering descriptors rather than a universal numerical score, allowing heterogeneous methods to be compared in terms of latency, observability, drift vulnerability, and deployment role.

Table 2. Distribution of included references and primary deployment-facing studies by evidence type, validation level, recording or modeling context, and study design.

Category	N	%
A. Included references by evidence type		
Primary deployment-facing studies	146	71.2
Reviews, consensus papers, or mechanistic background	40	19.5
Dataset or resource papers	11	5.4
Evaluation frameworks or reporting standards	5	2.4
Background or epidemiology	3	1.5
Total	205	100.0
B. Primary deployment-facing studies by validation level		
Offline or retrospective validation	84	57.5
Quasi-online or chronological validation	28	19.2
Prospective online validation	16	11.0
Chronic clinical deployment	18	12.3
C. Primary deployment-facing studies by dominant recording modality		
Invasive electrophysiology	40	27.4
Scalp EEG or other non-invasive EEG	22	15.1
Ambulatory, subcutaneous, wearable EEG	12	8.2
Chronic implant or device-based neuromodulation studies	33	22.6
Not primarily recording-based, mixed-modality, or not applicable	39	26.7
D. Primary deployment-facing studies by study design		
Retrospective	110	75.3
Prospective	36	24.7

Note: Percentages in Panel A use all 205 included references as the denominator; Panels B-D use the 146 primary deployment-facing studies. Invasive electrophysiology includes ECoG, SEEG, and mixed intracranial recordings. The ambulatory and wearable category includes subcutaneous EEG and wearable physiological monitoring used for deployment-facing seizure detection or forecasting. Percentages may not sum to 100 due to rounding.

Table 3. Worked examples of the four axis framework.

Example method	Temporal scale / latency	Observability / representation	Non-stationarity / drift	Deployment role
CHB-MIT deep learning seizure detector [18;19]	CHB-MIT-based deep learning seizure detection is typically evaluated on short EEG windows and can be implemented causally in principle, but deployment interpretation depends on event-level detection latency and continuous false alarm burden rather than segment-level accuracy alone.	The representation is sensor-level scEEG from a limited pediatric montage, typically using approximately 23 channels in the CHB-MIT configuration, with seizure annotations converted into segment-level labels and no explicit source or network-state model.	The binding limitation is distribution shift across patients, recordings, vigilance state, medication context, and recording sessions; operating-point drift appears clinically as changing false positives per unit time.	This supports retrospective or quasi-online continuous monitoring, but not real-time stimulation or chronic autonomous control without prospective latency, false alarm, and calibration evidence.
CCEP pipeline [20;21]	CCEP analysis operates at the millisecond scale; interpretation of the 0-20 ms early window depends on explicit artifact handling, response-window definition, and whether early neural signal retention is demonstrated.	The representation is a contact-level evoked response or protocol-conditioned effective-connectivity summary, rather than a direct whole-network readout; source ambiguity and incomplete sampling constrain interpretation.	The binding limitation is protocol and preprocessing dependence, including stimulation parameters, referencing, artifact suppression, response-window definition, and cross-session reproducibility.	This supports offline mapping, hypothesis generation, or periodic programming support, but not a millisecond-scale closed-loop control policy.
RNS reactive policy [22;23]	RNS operates with device-bounded sensing-to-stimulation timing, predefined detection pathways, and constrained stimulation actions. The fast loop is reactive, while parameter adaptation occurs on a slower clinical programming timescale.	The control variable is a detector-defined electrographic event observed through a limited implanted montage, not an independently measured epileptic-network state or whole-brain latent state.	The binding limitation is chronic detector drift, calibration decay, evolving electrographic patterns, and changing clinical response, managed through clinician-governed review and reprogramming over months.	This is real-time clinician-governed closed-loop control within a safety envelope, not unrestricted autonomous policy learning.

Note: The examples instantiate the same reportable descriptors in three representative settings: retrospective scEEG seizure detection, stimulation-evoked mapping, and implantable reactive stimulation. They are intended to show how deployment constraints can be made explicit for concrete method families rather than to provide exhaustive performance summaries or a universal numerical deployment score. scEEG, scalp EEG; CHB-MIT, Children's Hospital Boston-MIT; CCEP, cortico-cortical evoked potential; RNS, responsive neurostimulation.

Table 4. Validation levels for epilepsy neuromodulation algorithms.

Validation level	Supported inference	Core reporting and deployment metrics	Inference not supported
Offline validation	Whether a signal, feature, biomarker, or model can discriminate events, states, or outcomes under retrospective conditions.	Task-dependent discrimination metrics, including AUC, sensitivity, precision, F1 score, event-level detection rate, and retrospective false positives per unit time where applicable.	Does not establish prospective reliability, real-time feasibility, calibration stability, or acceptable clinical burden.
Quasi-online validation	Whether performance is preserved when data are evaluated in temporal order, with patient-specific calibration separated from later testing.	Sensitivity at a fixed false alarm burden, false positives per hour/day, calibration drift or decay, chronological split design, temporal leakage control, and performance on future held-out segments.	Does not establish behavior during live operation, clinician workload, response latency, or robustness to unforeseen recording and workflow conditions.
Prospective online validation	Whether the system can operate in real time on incoming data, with latency, warnings, uncertainty, and failures measured during live use.	Processing latency, prospective sensitivity, false alarm burden, warning horizon, time-in-warning, output stability, transition behavior, abstention rate, uncertainty calibration, and clinician review burden.	Does not by itself establish long term therapeutic value, chronic drift robustness, or safety under sustained adaptive use.
Chronic clinical deployment	Whether the system remains safe, interpretable, and clinically useful over prolonged use in realistic care settings.	Long term degradation curve, drift robustness, recalibration frequency, stability of operating point, clinician override rate, safety related failure modes, rollback/escalation events, and decision or treatment impact.	Does not remove the need for post-deployment monitoring, because patient state, treatment response, stimulation effects, and recording conditions may continue to evolve.

Note: Reporting and evaluation principles from EEG biomarker validation, seizure detection benchmarking, seizure forecasting evaluation, and calibration studies for clinical use are synthesized here. It distinguishes what each evaluation level can support from the stronger clinical inferences that require more prospective or longitudinal evidence. In practice, stronger deployment claims generally require the reporting and validation requirements of lower levels to have already been addressed. AUC, area under the ROC curve.

Table 5. Representative public epilepsy datasets and their clinical-use characteristics.

Dataset / resource	Sample size, duration, and sampling rate	Recording modality & annotation	Collection context and prospective/retrospective status	Deployment boundary / main limitation
Supervised in-hospital or retrospective clinical scEEG				
TUH EEG Corpus and TUSZ seizure subset [39;40]	TUH: 10,874 subjects, 16,986 sessions; TUSZ v1.2.0: 315 subjects, >504 h annotated; 250-512 Hz, majority 250 Hz	scEEG; 19 ch; seizure start/stop, type, channel-based and term-based labels	R archive; seizure subset enriched by triage and manual annotation	Large and clinically heterogeneous, but retrospective and selection-biased toward high-seizure-yield recordings; limited chronic realism
CHB-MIT Scalp EEG Database [18]	23 subjects; 664 segments; 198 seizures; typically 1-h segments; 256 Hz	scEEG; 18-23 ch; seizure onset/offset annotations per file	R clinical monitoring; segmented with short inter-file gaps	Small pediatric cohort under medication withdrawal; short per-subject horizon; does not represent ambulatory chronic monitoring
Harvard EEG Database (HEEDB) [41]	108,341 patients; 284,341 recordings; 2003-2024; sampling rate available from file metadata	Clinical scEEG; routine, EMU, and ICU; BIDS-compliant EDF; report-level labels and extracted EEG findings	R multi-hospital clinical archive; standardized on BDSP	Very large and clinically representative, but labels are partly report-level or weakly localized; does not substitute for prospective deployment validation
Long duration monitoring and chronic implanted sensing				
EPILEPSIAE [42]	275 patients; mean 165 h/patient; max 500 h; >40,000 h total; 250 Hz-2.5 kHz	Surface EEG (217 patients) and iEEG (58 patients); standardized seizure annotations including type, vigilance, semiology, and interictal events	R long term presurgical monitoring; commonly reused for quasi-prospective algorithm analysis	Rich temporal structure, but biased toward drug-resistant presurgical patients with high seizure yield; medication changes and access restrictions limit generalizability
NeuroVista chronic ECoG dataset [25;43]	9 subjects (Karoly analysis); mean 320 days/subject; 900 training + 2,879 testing days; 400 Hz	Chronic implanted ECoG; 16 intracranial contacts; automated seizure detections verified by investigators with audio and patient diaries	P chronic implanted data collection; commonly reused for retrospective or pseudo-prospective forecasting analysis	Strong chronic sensing realism, but not a fully closed-loop stimulation dataset; small, highly selected implanted cohort with device- and implantation-specific bias; small N limits population-level inference
Task-specific interictal event datasets				
Lin et al. IED dataset (2025) [44]	84 patients; 20 min/patient; 28 h total; 2,516 IED and 22,933 non-IED epochs; 500 Hz	scEEG; 19 ch (10-20) + T1/T2, A1/A2, ECG, EMG; IED type and wake/sleep annotated by ≥ 3 experts; 2 ms temporal resolution	R clinical EEG; designed for IED detection model training and simulated real-world testing	Spatial and consciousness-state labels available, but short-duration and single-center; not designed for continuous seizure detection or chronic deployment
Falach et al. iEEG sleep IED dataset (2024) [45]	25 patients; 1-5 min/patient; 76 min total; 857 iEEG channels; 852 annotated IED epochs; 2 kHz	Depth iEEG during overnight sleep; sub-second IED timing and brain-location labels per channel; inter-rater reliability assessed in 6 patients	R iEEG-BIDS dataset from 2 medical centers; clinical implantation with research sleep recordings	Rare expert-annotated intracranial sleep IED resource; but short-duration, sleep-specific, and regionally biased toward hippocampal/MTL recordings

Dataset / resource	Sample size, duration, and sampling rate	Recording modality & annotation	Collection context and prospective/retrospective status	Deployment boundary / main limitation
Omni-iEEG [46]	302 patients; 178 h total; 36,177 annotated HFO candidate events; original sampling rates retained, benchmark inputs resampled to 1000 Hz	Presurgical iEEG; ECoG and SEEG; expert labels for artifact, non-spkHFO, and spkHFO; SOZ, resection, and surgical outcome metadata	R harmonized multicenter public iEEG benchmark assembled from multiple open datasets	Large and standardized, but still based on retrospective presurgical recordings; HFO/spkHFO annotations remain task-specific and may not directly generalize to IED detection, seizure forecasting, or closed-loop deployment
Acute stimulation mapping and stimulation-evoked responses				
Parmigiani et al. simultaneous SEEG and hd-scEEG SPES resource [26]	36 patients; 323 artifact-free SPES sessions; simultaneous SEEG and 256-channel hd-EEG recorded at 1000 Hz	Simultaneous SEEG and hd-EEG during SPES; CCEP amplitude, latency, and spectral responses analyzed across physical, geometrical, and topological stimulation parameters	R presurgical evaluation cohort; open BIDS-formatted resource available through EBRAINS and OSF, with stimulating-contact locations, de-identified MRI, and digitized positions for 185 scEEG electrodes	Provides a cross-scale reference for linking intracranial stimulation responses to scEEG, but remains single-center, acute, protocol-specific, and constrained by clinical implantation geometry and limited stimulation sites per patient
PRIOS CCEP recordings [27]	7 patients aged 13-53 years; reported dataset-level recording duration not specified; 2048 Hz	iEEG/ECoG during SPES; explicit stimulation timing, CCEP/N1-peak labels, and electrode-position metadata; awake clinical and propofol conditions	R iEEG-BIDS dataset linked to the PRIOS study; protocol-specific recordings during awake clinical SPES and intraoperative SPES under propofol anesthesia	Small, protocol-specific acute stimulation-mapping dataset; limited reported recording duration and limited longitudinal context; should not be generalized to home monitoring, seizure forecasting, or autonomous closed-loop neuromodulation

Note: TUSZ, Temple University Hospital EEG Seizure Corpus; CHB-MIT, Children's Hospital Boston-MIT; HEEDB, Harvard Electroencephalography Database; IED, interictal epileptiform discharge; HFO, high frequency oscillation; SOZ, seizure onset zone; SPES, single-pulse electrical stimulation; scEEG, scalp EEG; hd-scEEG, high-density scalp EEG; iEEG, intracranial EEG; SEEG, stereo-electroencephalography; ECoG, electrocorticography; MTL, medial temporal lobe; BIDS, Brain Imaging Data Structure; OSF, Open Science Framework; R, retrospective; P, prospective. TUH, Temple University Hospital; CCEP, cortico-cortical evoked potential; EDF, European Data Format; EMU, epilepsy monitoring unit; ICU, intensive care unit; BDSP, Brain Data Science Platform; MRI, magnetic resonance imaging; hd-EEG, high-density EEG; spkHFO, spike-associated high frequency oscillation.

Table 6. Dataset-derived sources of methodological bias and deployment failure modes.

Dataset property / bias source	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 modeling and circadian/multidien analyses	Temporal leakage arises if splits ignore time; operational metrics should include FP/min, time-in-warning, and longitudinal stability
Cohort, device, and prevalence structure		
multicenter / 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; modeling 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: FC, functional connectivity; EC, effective connectivity; AUC, area under the ROC curve; FP/min, false positives per minute; PPV, positive predictive value; scEEG, scalp EEG; iEEG, intracranial EEG; SNR, signal-to-noise ratio.

Table 7. Selection framework for stimulation-artifact suppression method families.

Family	Best-fit scenarios	0-20 ms neural-signal retention	Online suitability	Minimum reporting requirements
Template / reconstruction	Stereotyped artifacts; stable SPES protocols; feasible template updating under limited drift	High	High; simple updates	Update rule; alignment; early-window residual metric; template drift or distortion check
Blanking / interpolation	Severe saturation or clipping; safety-driven mitigation; simple causal front-end	Low	Very high; minimal latency	Blanking window (ms); interpolation method; early samples removed; residual artifact estimate after recovery; statement of whether early responses were excluded from interpretation
Spectral / adaptive	Harmonic or narrowband artifacts, commonly in tACS/tES; separable spectral structure	Medium	Medium; filter- or adaptation-dependent	Filter causality or zero-phase use; bands; group delay or adaptation rule; residual artifact in the analyzed band or early window
BSS / projection / subspace	Multi-channel recordings with coherent artifact footprint; approximately stable artifact subspaces; sufficient channel count and stable spatial structure	Medium	Medium; channel-count dependent	Channels retained; subspace window; components removed or projected out; residual artifact and neural-signal loss assessment
Model / state-space	Non-stationary artifacts; protocol variation; explicit dynamics may improve tracking	High	Medium; often requires offline identification	Model form; identification procedure; early-window fit or residual error; residual artifact amplitude; online latency if used prospectively
Deep learning	Large and trustworthy training sets; complex artifacts; explicit protection against leakage and early-window distortion	Medium	Medium; inference latency validation required	Training design; label source; leakage controls; residual artifact and response-shape distortion assessment; cross-session or cross-protocol validation

Note: "0-20 ms neural-signal retention" refers to the potential to preserve interpretable signal structure in the early post-stimulation window used for early-response interpretation. The final column lists the minimum information needed to interpret early-response or perturbational effective connectivity claims. Ratings are qualitative, assume the best-fit scenarios listed for each method family, and are intended for method-family comparison rather than direct benchmarking across heterogeneous hardware platforms, stimulation protocols, channel counts, and software implementations. SPES, single-pulse electrical stimulation; CCEP, cortico-cortical evoked potential; tACS, transcranial alternating current stimulation; tES, transcranial electrical stimulation; BSS, blind source separation.

Table 8. Task framework for spontaneous-event sensing and SPES-evoked responses in neuromodulation.

Task	Output	Time scale / timing constraint	Evaluation / reporting metrics
IED detection	Event timing; optional subtype or morphology label	ms-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 s	Sensitivity versus false alarms/day, detection latency, and longitudinal stability under drift
Seizure prediction	Risk state / hazard estimate	minutes-days; warning horizon	Time-in-warning, false-warning burden, calibration, and prospective leakage-aware validation
SPES-evoked responses	Evoked-response summaries; optional ER/DR component characterization or CCEP-derived connectivity measures where responses are cortico-cortical	ms-hundreds of ms; 0-20 ms early-window fidelity where early components are interpreted	0-20 ms neural-signal retention where relevant, ER/DR component detectability, residual artifact reporting, and cross-pulse/cross-session consistency under protocol variation

Note: These tasks are highlighted as differing not only in output and time scale, but also in the form of evaluation required for clinically interpretable use. IED, interictal epileptiform discharge; FP/min, false positives per minute; SPES, single-pulse electrical stimulation; CCEP, cortico-cortical evoked potential; ER, early response; DR, delayed response.

Table 9. Connectivity framework for neuromodulation.

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 for implantation planning, pathway hypotheses, and stimulation-spread interpretation
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 for patient selection or target hypotheses
Scalp EEG/MEG	FC	Volume conduction, source leakage, inverse-model dependence, reference and montage effects; correlation, coherence, and phase-synchrony estimates can reflect common-reference or zero-lag confounds, whereas leakage-reduced variants may miss genuine near-zero-lag interactions; source-space analysis relocates rather than eliminates leakage	Monitoring and phenotyping: non-invasive state tracking and syndrome-level network characterization; invasive programming claims require independent validation
iEEG	FC / model-based or propagation-derived EC	Sparse and hypothesis-driven coverage; electrode density and sampling geometry affect apparent topology and hub structure; reference dependence; correlation, coherence, phase-synchrony, Granger-causality-like, and transfer-entropy estimates remain sensitive to sampling geometry, SNR, model order or estimator choice, and unobserved common inputs; not a valid whole-brain topology without adequate coverage	Patient-specific mapping, candidate contact ranking, and clinician-reviewed programming hypotheses when coverage, robustness, and clinical validation are reported
Generative or model-based inference from observed dynamics	EC	Model mis-specification; non-stationarity; parameter non-identifiability; sensitivity of inferred directionality to SNR, model structure, and observation assumptions	Prediction / programming support: directed propagation hypotheses interpreted within the EC boundary defined in the text and checked against longitudinal or clinical evidence
CCEP-derived perturbational connectivity	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, 0-20 ms neural-signal retention where relevant, and protocol robustness 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: SC, structural connectivity; FC, functional connectivity; EC, effective connectivity; DTI, diffusion tensor imaging; PET, positron emission tomography; fMRI, functional magnetic resonance imaging; MEG, magnetoencephalography; iEEG, intracranial EEG; CCEP, cortico-cortical evoked potential.

Table 10. Comparison of brain models and cross-modal transformations in epilepsy neuromodulation.

Model family	Representation / observability	Primary deployment role	Online feasibility	Fitting burden	Implementation burden	Dominant failure modes
Monitoring and control-design representations						
Latent-state models (HMM / state-space identification)	Sensor-level to low-dimensional latent state	Monitoring, triggering, risk estimation	Online-compatible with stable state definitions and recalibration rules	Threshold setting; periodic recalibration	Existing sensor stream plus state inference	State mislabeling, calibration drift, nonstationarity, limited interpretability
Reduced-order / surrogate control models (tracking control / iEEG-derived reduced dynamics)	Low-dimensional state or reduced dynamics derived from observed activity	Controller design, policy prototyping, constrained stimulation	Potentially online-compatible within fitted regimes	Patient-specific fitting to intracranial recordings	iEEG acquisition plus reduced-model fitting	Local validity, state drift, limited transfer outside fitted regime
Planning-oriented mechanistic and virtual models						
Mechanistic seizure models (Epileptor / neural-mass models)	Latent physiological variables and structured seizure dynamics	Mechanism clarification, hypothesis testing, EZN estimation	Generally offline or iterative	Priors, model fitting, posterior or sensitivity analysis	Simulation, optimization, or Bayesian inference infrastructure	Non-identifiability, parameter uncertainty, mismatch to routine clinical observables
Network ictogenicity / virtual resection (synchronizability / SIR / node ictogenicity)	Patient-specific network propensity measures	Candidate target ranking, presurgical planning	Offline simulation before intervention decisions	Connectome quality, electrode coverage, seizure-onset or EZN assumptions	MRI/DWI or electrode-localized iEEG/ECoG network construction	Connectome bias, incomplete sampling, model dependence, limited external validation
Personalized virtual-brain frameworks (VEP / virtual cohort)	Virtual patient-specific latent and network states with source-to-sensor mapping	In-silico stimulation testing, scenario exploration, synthetic benchmarking	Offline or iterative; not routine real-time	Multimodal fitting, source-to-sensor mapping, model inversion	Imaging-to-simulation pipeline with uncertainty propagation	Uncertainty accumulation, parameter degeneracy, incomplete sampling, simulation-to-outcome gap
Mapping and cross-modal transformation tools						
Electrical source imaging (ESI)	Sensor-to-source mapping	Localization support, presurgical mapping	Near-offline event review	Head model and inverse regularization choices	Head modeling, preprocessing, event selection	Head-model dependence, inverse ill-posedness, preprocessing and event-selection bias
Cross-modal / generative transformations (scalp-to-source / scalp-to-iEEG mapping)	Sensor-to-source or sensor-to-intracranial representations	Exploratory representation alignment, downstream validation	Training offline; routine real-time use not established	Paired datasets, cross-patient domain adaptation	Paired acquisition infrastructure plus decision-level validation	Domain shift, weak ground truth, leakage, decision-irrelevant reconstructions

Note: Rows are grouped according to the three roles used in the text: monitoring and control-design representations, planning-oriented mechanistic and virtual models, and mapping or cross-modal transformation tools. Representation, deployment role, feasibility, fitting burden, implementation burden, and dominant failure modes are compared rapidly here; detailed interpretation is provided in the corresponding subsections. HMM,

hidden Markov model; EZN, epileptogenic zone network; VEP, Virtual Epileptic Patient. iEEG, intracranial EEG; MRI, magnetic resonance imaging; DWI, diffusion-weighted imaging; ECoG, electrocorticography; SIR, susceptible-infected-recovered model; ESI, electrical source imaging.

Table 11. Policy interface for closed-loop neuromodulation.

Policy component	Minimum reporting / governance requirement
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 balance, charge density, stimulation duration or 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 / update	Governed adaptation mechanisms: clinician-in-the-loop reprogramming or, where justified, bounded autonomous adaptation with predefined update cadence or triggering conditions, supported by auditability, logging, versioning, rollback, and criteria for returning to a prior clinician-approved setting.

Note: Closed-loop systems are highlighted as needing to be understood not simply as detectors linked to stimulators, but as governed control policies defined by sensed inputs, inferred states, bounded actions, safety constraints, fallback logic, and explicit update mechanisms. These components also provide a minimum reporting checklist for evaluating whether an adaptive or closed-loop system is clinically governable rather than only technically functional. IED, interictal epileptiform discharge; EC, effective connectivity.

Table 12. Representative clinical stimulation parameters and programmable constraints in epilepsy neuromodulation.

Therapy / device class	Control logic	Representative clinical parameters	Duty ratio / update cadence	Programmable levers, supervision, and safety envelope
RNS	Intracranial responsive stimulation triggered by detected electrographic activity.	Common pivotal-trial responsive settings include two bursts at 100-200 Hz, 160 µs pulse width, 100 ms burst duration, charge density of 4-5 µC/cm^2, and current of 4.5 mA. NeuroPace clinical guidance recommends initial settings of 200 Hz, 160 µs pulse width, 100 ms burst duration, and 1.0 mA current; broader parameter ranges and alternative stimulation strategies should be interpreted separately. [22;187]	Not defined by a fixed cyclic ON/OFF schedule; therapy exposure is better described by detections or therapies delivered, bursts per detection, burst duration, charge density, current amplitude, and follow-up titration, with programming adjustments typically occurring during clinician follow-up over weeks to months rather than through continuous autonomous online updates.	Detection settings, sensing contacts, stimulation pathway, current, frequency, pulse width, burst duration, charge density, burst configuration, and tolerability limits are adjusted longitudinally under clinician supervision.
VNS	Intermittent peripheral neuromodulation, with optional magnet-triggered or cardiac-responsive features in some systems.	Common conventional programming uses intermittent stimulation, typically described as 30 s ON every 5 min or 30 s ON / 5 min OFF, with output current, pulse width, frequency, and duty cycle adjusted according to tolerability and response. [173;185;186]	A 30 s ON / 5 min OFF schedule corresponds to an approximate 9-10% duty cycle, depending on whether the denominator is expressed as the OFF interval alone or the full ON-plus-OFF cycle; duty cycle can be increased by shortening OFF time.	Output current, pulse width, frequency, ON/OFF timing, magnet mode or cardiac-responsive triggering, and side-effect management define the practical programming surface.
ANT-DBS	Usually cyclic open-loop DBS of a thalamic network target; sensing-enabled or adaptive variants remain investigational in epilepsy.	Representative SANTE-style settings use high frequency stimulation at 145 Hz, 90 µs pulse width, 5 V amplitude, and cycling stimulation of 1 min ON / 5 min OFF; clinical practice may deviate from pivotal-trial defaults. [182–184;187]	A 1 min ON / 5 min OFF schedule corresponds to an approximate 16.7% duty ratio.	Amplitude or current, pulse width, frequency, contact configuration, cycling schedule, target selection, and side-effect mitigation are individualized over follow-up.

Note: Values are representative clinical settings or commonly reported programming paradigms rather than universal prescriptions. Device manuals or regulatory materials should be cited for full programmable ranges and hard safety limits, whereas peer-reviewed clinical studies or reviews should be cited for commonly used programming paradigms. Charge density refers to pulse-level local electrical exposure, whereas duty ratio refers to cumulative temporal exposure; for responsive stimulation, therapy exposure is often better described by detections or therapies delivered, burst duration, charge density, current amplitude, and follow-up titration rather than by a fixed ON/OFF cycle. Clinically reported operating regimes and programmable constraints are summarized here, not to replace device-specific manuals, regulatory labeling, or patient-specific safety programming. RNS, responsive neurostimulation; VNS, vagus nerve stimulation; ANT-DBS, anterior nucleus of the thalamus deep brain stimulation; DBS, deep brain stimulation.

Table S1. Structured scoping search strategy, search window, and database-specific queries used for record identification.

Search component	Scope	Search terms and implementation
Databases searched	Bibliographic and interdisciplinary databases	PubMed, IEEE Xplore, ScienceDirect, and Google Scholar. PubMed and IEEE Xplore were used as primary bibliographic databases. ScienceDirect and Google Scholar were used as broad supplementary discovery sources; only the records actually screened in relevance-ranked order were counted in the flow diagram.
Search period	Main search window	January 1, 1990 to June 2, 2026. Ahead-of-print and in-press articles available by 2 June 2026 were eligible.
Block 1: epilepsy and neuromodulation	Disease, clinical target, and intervention terms	epilepsy; seizure; epileptiform activity; interictal epileptiform discharge (IED); neuromodulation; stimulation; responsive neurostimulation (RNS); deep brain stimulation (DBS); vagus nerve stimulation (VNS); closed-loop stimulation; adaptive stimulation.
Block 2: recording and sensing modalities	Signal acquisition and monitoring terms	scalp electroencephalography (scEEG); electroencephalography (EEG); intracranial electroencephalography (iEEG); electrocorticography (ECoG); stereo-electroencephalography (SEEG); local field potentials (LFPs); chronic EEG; ambulatory EEG; stimulation-while-sensing recordings.
Block 3: methodological and deployment targets	Algorithmic, modeling, and translational terms	artifact handling; artifact removal; artifact suppression; event detection; IED detection; seizure detection; seizure forecasting; seizure prediction; evoked-response quantification; single-pulse electrical stimulation (SPES); cortico-cortical evoked potentials (CCEPs); source localization; connectivity analysis; functional connectivity; effective connectivity; brain modeling; brain modelling; virtual brain; electrical source imaging; machine learning (ML); deep learning (DL); adaptive control; calibration; drift; non-stationarity; clinical deployment.
PubMed database query	Biomedical and clinical query	(("epilepsy"[Title/Abstract] OR "epileptic"[Title/Abstract] OR "epileptiform"[Title/Abstract] OR "interictal epileptiform discharge"[Title/Abstract] OR "interictal epileptiform discharges"[Title/Abstract] OR "focal epilepsy"[Title/Abstract]) AND ("EEG"[Title/Abstract] OR "electroencephalography"[Title/Abstract] OR "scalp EEG"[Title/Abstract] OR "intracranial EEG"[Title/Abstract] OR "iEEG"[Title/Abstract] OR "ECoG"[Title/Abstract] OR "SEEG"[Title/Abstract] OR "local field potential"[Title/Abstract] OR "local field potentials"[Title/Abstract]) AND ("artifact removal"[Title/Abstract] OR "artifact suppression"[Title/Abstract] OR "stimulation artifact"[Title/Abstract] OR "IED detection"[Title/Abstract] OR "interictal epileptiform discharge detection"[Title/Abstract] OR "seizure detection"[Title/Abstract] OR "seizure prediction"[Title/Abstract] OR "seizure forecasting"[Title/Abstract] OR "single-pulse electrical stimulation"[Title/Abstract] OR "SPES"[Title/Abstract] OR "cortico-cortical evoked potential"[Title/Abstract] OR "cortico-cortical evoked potentials"[Title/Abstract] OR "CCEP"[Title/Abstract] OR "CCEPs"[Title/Abstract] OR "seizure onset zone"[Title/Abstract] OR "electrical source imaging"[Title/Abstract] OR "brain modeling"[Title/Abstract] OR "brain modelling"[Title/Abstract] OR "virtual brain"[Title/Abstract] OR "responsive neurostimulation"[Title/Abstract] OR "closed-loop stimulation"[Title/Abstract] OR "adaptive stimulation"[Title/Abstract] OR "adaptive control"[Title/Abstract])) AND ("1990/01/01"[Date - Publication] : "2026/06/02"[Date - Publication]) NOT (animals[mh] NOT humans[mh]). This query identified 2,721 records before deduplication.
IEEE Xplore database query	Engineering and signal-processing query	("All Metadata":"epilepsy" OR "All Metadata":"epileptic seizure" OR "All Metadata":"epileptiform" OR "All Metadata":"interictal epileptiform discharge") AND ("All Metadata":"EEG" OR "All Metadata":"intracranial EEG" OR "All Metadata":"iEEG" OR "All Metadata":"ECoG" OR "All Metadata":"SEEG" OR "All Metadata":"local field potential" OR "All Metadata":"LFP") AND ("All Metadata":"artifact removal" OR "All Metadata":"artifact suppression" OR "All Metadata":"stimulation artifact" OR "All Metadata":"IED detection" OR "All Metadata":"interictal epileptiform discharge detection" OR "All Metadata":"EEG seizure detection" OR "All Metadata":"epileptic seizure detection" OR "All Metadata":"real-time seizure detection" OR "All Metadata":"patient-specific seizure detection" OR "All Metadata":"wearable seizure detection" OR "All Metadata":"seizure forecasting" OR "All Metadata":"seizure prediction" OR "All Metadata":"single-pulse electrical stimulation" OR

Search component	Scope	Search terms and implementation
		"All Metadata":"SPES" OR "All Metadata":"cortico-cortical evoked potential" OR "All Metadata":"CCEP" OR "All Metadata":"electrical source imaging" OR "All Metadata":"effective connectivity" OR "All Metadata":"closed-loop stimulation" OR "All Metadata":"adaptive stimulation" OR "All Metadata":"adaptive control" OR "All Metadata":"responsive neurostimulation"). The publication-year filter was set to 1990-2026; records published after June 2, 2026, if any, were removed during screening. This query identified 1,572 records before de-duplication.
ScienceDirect database query	Thematically partitioned sub-queries, broad biomedical-engineering coverage	Because ScienceDirect's advanced search interface restricts each query to a maximum of eight Boolean connectors per field, the full search logic was implemented as four thematically partitioned sub-queries, each combining core epilepsy and EEG/iEEG terms with one methodological sub-domain: (i) detection/forecasting: ("epilepsy" OR "seizure") AND ("EEG" OR "iEEG" OR "ECoG") AND ("seizure detection" OR "seizure forecasting" OR "seizure prediction") (4,024 records); (ii) stimulation/neuromodulation: ("epilepsy" OR "seizure") AND ("EEG" OR "iEEG" OR "SEEG") AND ("SPES" OR "CCEP" OR "closed-loop neuromodulation" OR "adaptive stimulation") (1,654 records); (iii) connectivity/modeling, restricted to intracranial recording modalities to align with the deployment-facing focus of this review: ("epilepsy" OR "seizure") AND ("iEEG" OR "ECoG" OR "SEEG") AND ("functional connectivity" OR "effective connectivity" OR "network analysis") (1,300 records); and (iv) artifact handling/clinical deployment: ("epilepsy" OR "seizure") AND ("EEG" OR "iEEG") AND ("artifact suppression" OR "artifact removal" OR "stimulation artifact" OR "clinical deployment") (1,697 records). The publication-year filter was set to 1990-2026 for all sub-queries; records published after June 2, 2026, if any, were removed during screening. These sub-queries returned 8,675 records in total. ScienceDirect was treated as a supplementary, broad-coverage source rather than an exhaustively screened primary database. Records from each sub-query were sorted by relevance, and the first 100 records per sub-query were screened at the title level, yielding 400 ScienceDirect records taken forward for de-duplication.
Google Scholar supplementary search	Citation tracking and targeted supplementary search	Google Scholar was used for citation tracking, supplementary cross-checking, and targeted gap searches. Six targeted phrase searches were used: "responsive neurostimulation" epilepsy long term outcome (4,500 results), "seizure forecasting" long term EEG (1,380 results), "single-pulse electrical stimulation" "cortico-cortical evoked potential" epilepsy (402 results), "stimulation artifact" suppression EEG (1,550 results), "effective connectivity" epilepsy stimulation (10,700 results), and "adaptive neuromodulation" epilepsy (362 results). Because Google Scholar does not provide a stable or reproducible exhaustive export, it was treated as a supplementary discovery source rather than an exhaustively screened primary database. The first 100 relevance-ranked records per phrase search were screened at the title level, yielding 600 Google Scholar records taken forward for de-duplication.
Revision-stage update	Targeted update for reviewer-identified gaps	Targeted searches were added during revision for prospective validation, chronic and ambulatory iEEG, stimulation-while-sensing datasets, early-window artifact recovery, connectivity estimator limitations, long term drift and calibration, and adaptive neuromodulation. Final update: June 2, 2026.

Note: IED, interictal epileptiform discharge; EEG, electroencephalography; iEEG, intracranial EEG; ECoG, electrocorticography; SEEG, stereo-electroencephalography; LFP, local field potential; SPES, single-pulse electrical stimulation; CCEP, cortico-cortical evoked potential; RNS, responsive neurostimulation; DBS, deep brain stimulation; VNS, vagus nerve stimulation; ML, machine learning; DL, deep learning; scEEG, scalp EEG.