Review Article ISSN 2693-2490

Journal of Psychology and Neuroscience

Machine Learning and Brain-Computer Interfaces for Medical and Psychiatric Translation, 2024-2026

Dr. Chuck Easttom

Vanderbilt and Georgetown University, United States.

***Corresponding Author**
Dr. Chuck Easttom,
Vanderbilt and Georgetown University,
United States.

Submitted: 13 Jul 2026; **Accepted:** 25 Jul 2026; **Published:** 30 Jul 2026

Citation: Easttom, C. (2026). Machine Learning and Brain-Computer Interfaces for Medical and Psychiatric Translation, 2024-2026. *J Psychol Neurosci; 8*(4):1-11. DOI : https://doi.org/10.47485/2693-2490.1172

Abstract

Brain-computer interfaces (BCIs) are becoming closed-loop medical cyber-physical systems in which neural activity is decoded, transformed into a control or therapeutic signal, and returned to the patient through speech synthesis, prosthetic actuation, neurofeedback, electrical stimulation, or clinician-supervised rehabilitation. This survey synthesizes machine-learning trends from January 2024 through May 2026 with emphasis on medical and psychiatric translation. The dominant methodological transition is from small supervised decoders toward multimodal representation learning, self-supervised pretraining, domain adaptation, latent-dynamics stabilization, uncertainty-aware control, and clinically constrained online learning. The biomedical use cases are heterogeneous: non-invasive EEG and fNIRS support neurorehabilitation, cognitive-state monitoring, depression phenotyping, and neurofeedback; ECoG, sEEG, and intracortical arrays support high-performance communication and dexterous motor neuroprostheses; closed-loop DBS and responsive neurostimulation motivate psychiatric biomarkers for treatment-resistant depression, obsessive-compulsive disorder, post-traumatic stress disorder, addiction, and affective dysregulation. A unifying technical problem is generalization under severe distribution shift caused by subject variability, electrode impedance, medication state, vigilance, neuropsychiatric symptom fluctuation, and tissue-electrode nonstationarity.

Index Terms: Brain-computer interface, EEG, ECoG, intracortical BCI, psychiatric neuromodulation, depression, obsessive-compulsive disorder, PTSD, neurofeedback, speech neuroprosthesis, foundation model, self-supervised learning, adaptive decoding, closed-loop control, privacy-preserving machine learning.

Keywords: Foundation Models, Adaptive Decoding, Neuroprostheses, Closed-Loop Therapeutics, and Neural Data Governance.

Introduction and Clinical Scope

The original manuscript framed BCIs as statistical control systems that convert neural observations into communication, control, stimulation, or therapeutic feedback. The clinical version of the problem is stricter: the decoder is not an isolated classifier but a safety-critical component embedded in a patient-specific loop containing neural tissue, symptoms, assistive hardware, the caregiver, and the clinical workflow. A model that maximizes offline accuracy but requires prolonged supervised calibration, exhibits unstable latency, or fails silently during medication or arousal changes has limited translational value even when its benchmark score is high (Ding et al., 2026; Saibene et al., 2024; Hameed et al., 2025; Elashmawi et al., 2024; Williams et al., 2025).

The 2024-2026 literature shows rapid progress in foundation models, transformer-CNN hybrids, online speech synthesis, continuous high-degree-of-freedom control, latent alignment, and wearable multimodal sensing. However, the translational bottleneck remains the same across neurology and psychiatry: neural data are sparse, nonstationary, patient-specific, expensive to label, and sensitive. In psychiatric applications, additional complications arise because the latent variable of interest may be a dimensional symptom state such as anhedonia, threat reactivity, craving, hallucination burden, impulsivity, affective arousal, or intrusive-thought probability rather than an externally observable movement label.

For this reason, this survey treats BCI machine learning as a joint estimation-control-governance problem. The technical stack must support representation learning, uncertainty quantification, continual calibration, causal or quasi-causal biomarker validation, algorithmic fairness, and privacy-preserving deployment. The medical emphasis is not limited to motor restoration; it includes speech restoration in ALS and stroke, stroke rehabilitation, epilepsy monitoring, ICU neurophysiology, adaptive DBS, closed-loop psychiatric neuromodulation, and neurofeedback for affective and attentional disorders.

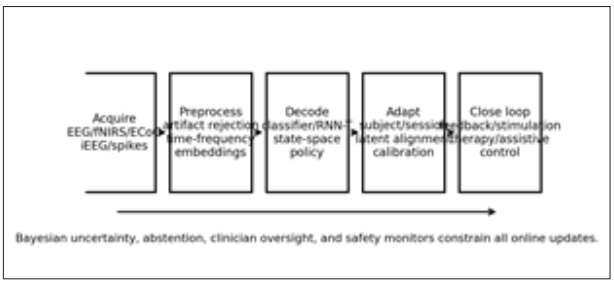

Figure 1: Closed-loop BCI machine-learning pipeline for medical and psychiatric systems. The model is embedded within a feedback loop and therefore must be evaluated as a dynamical clinical system rather than as a static offline decoder.

Mathematical Formulation of Translational BCI Learning

Let X_t in R^{C x T} denote a multichannel neural window with C sensors and T samples. The decoding target y_t may be a motor class, cursor velocity, phoneme sequence, psychiatric symptom state, stimulation command, or latent clinical endpoint. A minimal encoder-decoder representation is

$$z_t = \phi_\theta\{A(X_{t-k:t}, m_t, q_t)\}$$

$$\hat{y}_t = g_\psi(z_t, c_t) \quad (1)$$

where A is preprocessing and artifact handling, m_t denotes montage or channel metadata, q_t denotes data-quality variables such as impedance or motion state, z_t is a latent neural representation, and C_t is clinical context such as medication state, task, symptom scale, or stimulation history.

A clinically realistic objective cannot be reduced to a single cross-entropy term. A more informative risk functional is

$$\min_{\theta,\psi} \frac{1}{N}\sum_i \mathcal{L}_{\text{task}}\{g_\psi[\phi_\theta(A(X_i))], y_i\} + \lambda D_{\text{dom}}(z_i, s_i) + \beta\mathcal{L}_{\text{SSL}}(X_i) + \eta R_{\text{safe}} + \kappa R_{\text{priv}} \quad (2)$$

where D_{dom} penalizes subject, session, site, or device shift; L_{SSL} is a self-supervised objective such as masked reconstruction or contrastive prediction; R_{safe} penalizes unsafe commands, delayed feedback, or stimulation excursions; and R_{priv} captures privacy leakage from model parameters or shared gradients. The coefficients are not merely hyperparameters; they encode clinical priorities. For example, a depression-monitoring wearable may weight privacy and false-positive symptom escalation more strongly than raw classification accuracy, whereas a speech neuroprosthesis may weight real-time latency and false phrase emission.

For continuous neuroprosthetic control, the problem is better represented as state estimation. A linear-Gaussian baseline remains useful because it reveals the information pathway:

$$x_{t+1} = Fx_t + Bu_t + w_t,$$
$$n_t = Hx_t + v_t,$$
$$u_t = \pi_\phi\{p(x_t \mid n_{1:t})\} \quad (3)$$

where x_t is latent intent or symptom state, n_t is neural activity, u_t is the actuator or stimulation command, and pi_{phi} is the policy. Modern recurrent neural networks, sequence transducers, diffusion models, and transformer decoders can be interpreted as nonlinear generalizations of this estimator-control loop.

For EEG and MEG covariance methods, the Riemannian distance between two symmetric positive-definite covariance matrices C_1 and C_2 remains an important low-data baseline:

$$\delta_R(C_1, C_2) = \left\|\log(C_1^{-1/2} C_2 C_1^{-1/2})\right\|_F \quad (4)$$

Its persistence as a strong benchmark is a reminder that domain-specific inductive bias can outperform scale when labeled neural data are scarce (Chevallier et al., 2024), (Schirrmeister et al., 2017; Lawhern et al., 2018; Lotte et al., 2018).

Foundation Models and Self-Supervised Neural Representations

The most visible algorithmic shift is the attempt to construct reusable encoders for neural time series. LaBraM tokenizes EEG into channel patches and uses vector-quantized neural spectrum prediction; it reports pretraining on approximately 2,500 h of EEG from about 20 datasets (Jiang et al., 2014). NeuroLM increases the scale and language-model analogy by treating EEG tokens as a foreign-language stream, using temporal-frequency tokenization, autoregressive modeling, instruction tuning, and a reported approximately 25,000 h EEG corpus for its largest variant (Jiang et al., 2025).

The clinical motivation is straightforward: medical and psychiatric BCI tasks cannot be labeled at web scale. A psychiatric patient cannot be asked to provide thousands of precisely labeled episodes of intrusive thought, craving, hallucination, or suicidal ideation; an ALS speech-neuroprosthesis user cannot provide unlimited clean ground-truth articulation; and a stroke patient changes physiologically during rehabilitation. Self-supervised learning therefore attempts to harvest structure from unlabeled data before task-specific fine-tuning.

Table 1: Medical and Psychiatric BCI Task Taxonomy

Application	Clinical target	Primary ML/clinical endpoints
Communication restoration	ALS, brainstem stroke, locked-in syndrome, anarthria; output is text, acoustics, or personalized voice.	WER, latency, semantic error, user correction burden, privacy of intended speech.
Motor restoration	Cervical spinal cord injury, stroke, limb loss; output is cursor, hand trajectory, prosthesis, exoskeleton, or FES.	Target acquisition rate, DOF, closed-loop stability, fatigue, safe stop and collision risk.
Neurorehabilitation	Stroke, traumatic brain injury, cerebral palsy; model estimates motor intent or effort and adapts feedback.	Longitudinal functional outcomes, therapist integration, decoder-user co-adaptation.
Epilepsy/ICU monitoring	Seizure detection, brain-state deterioration, sedation and coma monitoring.	False alarm burden, detection delay, calibration across pathology and medications.
Psychiatric monitoring	MDD, bipolar depression, anxiety, PTSD, ADHD, psychosis; model estimates dimensional symptom biomarkers.	Label validity, fairness, ecological validity, false escalation, privacy and stigma.
Closed-loop psychiatry	Treatment-resistant depression, OCD, PTSD, addiction; stimulation or feedback is controlled by biomarkers.	Causal biomarker evidence, stimulation safety, symptom-state drift, autonomy, agency.

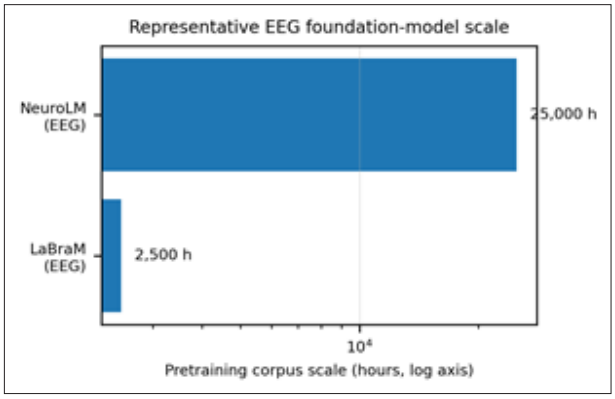

Figure 2: Representative pretraining-corpus scale for EEG foundation models reported in recent literature. The logarithmic axis emphasizes that model scale is increasing faster than standardized clinical validation.

Current evidence is conditional rather than decisive. Reviews and benchmarks report heterogeneous protocols, sensitivity to preprocessing, and weak linear-probe performance in some settings (Kuruppu et al., 2025; Dai et al., 2026; Yang et al., 2026), (Xiong et al., 2025), (Portmann et al., 2026). DeeperBrain and similar neuro-grounded models respond to this limitation by embedding volume-conduction, neurodynamical, and spectral priors into the architecture and objective (Wang et al., 2026). For medical use, the key question is not whether a foundation model has more parameters than an EEGNet-like decoder; it is whether it improves calibration-constrained, leakage-free, patient-independent generalization under prospectively specified clinical endpoints.

Self-supervised Objectives

Masked time-frequency reconstruction, contrastive predictive coding, autoregressive next-token prediction, and cross-modal alignment can be written as special cases of density-ratio or reconstruction learning. A contrastive neural objective is

$$\mathcal{L}_{\mathrm{NCE}} = -\log \frac{\exp[\mathrm{sim}(z_t, z_t^+)/\tau]}{\sum_{j}^{K} \exp[\mathrm{sim}(z_t, z_j)/\tau]} \quad (5)$$

Where z_{t+} is a temporally or semantically positive sample and z_j are negatives. In psychiatric phenotyping, negative sampling must be clinically controlled: two segments from different symptom states may be physiologically similar, and two segments from the same diagnostic category may be clinically heterogeneous.

Domain Adaptation and Calibration

Cross-subject adaptation is essential because scalp geometry, electrode impedance, skull conductivity, cortical anatomy, medications, sleep, psychiatric state, and task engagement induce domain shift. Maximum mean discrepancy provides one widely used alignment penalty:

$$D^2_{\mathrm{MMD}}(P_s, P_t) = \left\| \mathbb{E}_{z\sim P_s}[\Phi(z)] - \mathbb{E}_{z\sim P_t}[\Phi(z)] \right\|^2_{\mathcal{H}} \quad (6)$$

However, a clinical evaluation must distinguish unsupervised target adaptation from supervised target calibration. The amount of target data, labels, and clinician time should be reported as an outcome variable, not hidden as an implementation detail (Liang et al., 2024; Otarbay et al., 2025; Wu, 2025).

Table 2: Self-Supervised Objectives for Clinical BCI Data

Objective	Technical mechanism	Clinical implication
Masked reconstruction	Temporal segments, frequency patches, or channels are masked and reconstructed.	Robust pretraining with unlabeled EEG; risk of learning artifact statistics.
Contrastive prediction	Positive pairs may share time, task, subject, or modality; negatives approximate alternative states.	Sensitive to psychiatric label ambiguity and false negative sampling.
Autoregressive token modeling	Discrete neural tokens are predicted from prior context.	Useful for speech and continuous monitoring; latency and hallucination control are critical.
Cross-modal alignment	EEG/fNIRS/EMG/EOG/audio/text are embedded in a shared representation.	Supports rehabilitation and affective monitoring; confounding can be amplified.
Neurodynamics priors	Spectral power, functional connectivity, cross-frequency coupling, and complexity constrain learning.	Improves interpretability and may reduce sample requirements in clinical cohorts.

Architectures: Hybrid Inductive Bias Rather Than Generic Attention

Recent BCI architectures rarely use a generic transformer alone. EEG and fNIRS contain local spectral dynamics, spatial montage structure, ocular and muscular artifacts, nonuniform sampling, and task-locked transients. ECoG and intracortical recordings add spatially local high-gamma or spiking structure and small-cohort constraints. Consequently, successful systems typically combine convolutional front ends, temporal convolutional networks, attention blocks, recurrent state-space models, graph encodings, or tokenizers that preserve channel geometry (Liao & Wang, 2024; Klein et al., 2025; Altaheri et al., 2025; Vafaei & Hosseini, 2025).

A clinically deployable architecture must satisfy four simultaneous constraints: physiological validity, sample efficiency, real-time operation, and failure transparency. In psychiatric applications, interpretability is especially important because the model may influence a stimulation decision or symptom escalation. A black-box affect classifier trained on small convenience cohorts is not clinically equivalent to a validated biomarker.

For sequence BCIs, decoding is often framed as a conditional sequence model:

$$p(y_{1:L} \mid X_{1:T}) = \prod_{\ell=1}^{L} p(y_\ell \mid y_{<\ell}, h_{1:T}), \quad h_t = f_\theta(X_{1:t}) \tag{7}$$

Connectionist temporal classification, RNN transducers, attention-based encoder-decoders, and language-model rescoring instantiate different alignment assumptions. In speech neuroprostheses, the alignment is neurophysiological rather than purely acoustic: attempted articulation, inner speech, phonemes, prosody, and acoustic synthesis may not share identical neural timing.

Medical Neuroprostheses: Speech, Voice, and High-Dof Motor Control

Speech and voice neuroprostheses are the strongest recent examples of ML-driven BCI translation. They transform BCI from trial-level classification into streaming sequence transduction with explicit latency, language modeling, acoustic synthesis, and user-feedback loops. Online speech synthesis from a chronic BCI in ALS, streaming brain-to-voice synthesis with 80-ms neural decoding increments, and instantaneous voice synthesis from 256 intracortical microelectrodes demonstrate that clinically meaningful communication may require personalized acoustic output rather than text alone (Angrick et al., 2024; Littlejohn et al., 2025; Wairagkar et al., 2025; Card et al., 2024; Willett et al., 2023; Metzger et al., 2023).

The signal-processing burden is substantial. The decoder must distinguish silence from intended speech, map cortical activity to phonemic or articulatory trajectories, constrain output with language priors, preserve personal voice identity where consent permits, and avoid unauthorized inference of inner speech. For a patient with severe dysarthria or anarthria, the training signal may be weak or absent, requiring surrogate labels, alignment models, or model-based synthesis (Littlejohn et al., 2025), (Wairagkar et al., 2025).

High-degree-of-freedom motor control imposes a different set of constraints. Finger-level decoding and quadcopter game control demonstrated continuous control of three independent finger groups with thumb motion in two dimensions, yielding four degrees of freedom, 76 targets per minute, and a mean completion time of 1.58 +/- 0.06 s in a human participant with tetraplegia (Willsey et al., 2025). This result is clinically important not only because of the metric, but because recreation and social connection are patient-prioritized outcomes. Medical BCIs should therefore optimize quality-of-life utility, not only laboratory throughput.

Information transfer rate remains a useful but incomplete communication metric:

$$\mathrm{ITR} = \frac{60}{T}\left[\log_2 N + P\log_2 P + (1-P)\log_2\{(1-P)/(N-1)\}\right] \ \text{bits/min} \tag{8}$$

Here N is the number of commands, P is accuracy, and T is trial duration in seconds. The metric is inadequate when errors have

unequal clinical cost: an unintended letter, a wrong word, and a dangerous wheelchair command are not equivalent.

Table 3: Representative Medical Neuroprosthesis Demonstrations

System class	Neural/ML substrate	Translational contribution
Online speech synthesis	Chronically implanted BCI in ALS; neural-to-speech output.	Streaming synthesis, personalization, latency, clinical speech restoration (Angrick et al., 2024).
Streaming brain-to-voice	High-density surface recordings; RNN-transducer style synthesis in 80-ms increments.	Naturalistic communication with continuous decoding and personalized voice (Littlejohn et al., 2025).
Instantaneous voice synthesis	256 microelectrodes in ventral precentral gyrus; closed-loop audio feedback.	Decodes phonemic and paralinguistic control such as intonation (Wairagkar et al., 2025).
Rapid speech calibration	High-performance spelling or speech decoding from implanted arrays.	Rapid calibration, language-model constraints, and text output (Card et al., 2024; Willett et al., 2023; Metzger et al., 2023).
Finger/quadcopter BCI	Four-DOF finger decoding in tetraplegia; virtual quadcopter control.	76 targets/min, 1.58 +/- 0.06 s completion, patient-prioritized leisure (Willsey et al., 2025).
Latent stabilization	NoMAD aligns recurrent latent dynamics without supervised recalibration.	Weeks-to-months stability; reduced clinical recalibration burden (Karpowicz et al., 2025).

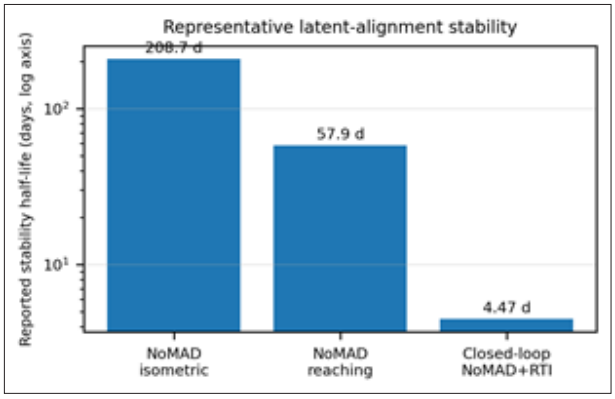

Figure 3: Representative half-life values reported for latent-dynamics stabilization with NoMAD. Values are plotted on a logarithmic axis because clinical recalibration burden is multiplicative over long-term use.

Psychiatric Applications: From Biomarkers to Closed-Loop Therapeutics

Psychiatric BCI applications should be treated as a spectrum rather than a single category. At one end are passive diagnostic or monitoring systems, such as EEG-based depression detection, workload estimation, or affective-state monitoring. At the other end are active closed-loop therapies, such as adaptive DBS, responsive neurostimulation, and neurofeedback systems that deliver feedback or stimulation contingent on decoded brain state. Between these extremes are hybrid systems that combine wearable physiology, behavioral context, ecological momentary assessment, and neural data to infer patient-specific symptom trajectories.

Major depressive disorder (MDD) is a paradigmatic case. A 2025 systematic review of EEG-based machine-intelligence methods for depression diagnosis and monitoring screened 938 candidate studies and retained 139 under PRISMA-style criteria (Nassibi et al., 2025). This scale indicates methodological growth, but not clinical readiness. Psychiatric labels are often subjective, comorbid, medication-dependent, and temporally unstable. A model trained to separate patients from controls at rest may not predict treatment response, relapse, suicidality, or function. Therefore, depression BCIs should prioritize prognostic validity, calibration across medications, fairness, and prospective validation.

Obsessive-compulsive disorder, PTSD, addiction, and severe anxiety are natural candidates for closed-loop neuromodulation because symptoms may be episodic and circuit-based. However, closed-loop psychiatry requires a biomarker that is both predictive and controllable. A biomarker can correlate with symptoms but still be unsuitable for stimulation control if it is epiphenomenal, confounded by motion or medication, or insensitive to therapeutic intervention. The relevant control-theoretic target is not classification accuracy alone but closed-loop symptom regulation under safety constraints:

$$u_t^* = \arg\min_{u \in U} \mathbb{E}[\ \ell(s_{t+1}, s_{\text{ref}}) + \rho \|u - u_{t-1}\|_2^2\ ] \quad \text{s.t. } \Pr(\text{harm} \mid u, z_t) \le \varepsilon \tag{9}$$

Here s_t is the latent symptom state, u_t is feedback or stimulation, rho penalizes abrupt dose changes, and epsilon is a clinician-specified upper bound on risk. This formalism is applicable to adaptive DBS, exposure-therapy augmentation, neurofeedback, and wearable-triggered interventions.

For ADHD and attentional disorders, neurofeedback and passive BCI overlap. EEG rhythms can provide feedback about attentional state, but the clinically relevant outcome is not merely modulation of theta/beta ratios or classifier accuracy; it is sustained improvement in function, school performance, impulsivity, and adverse-effect profile compared with standard treatment. Similar caution applies to PTSD and anxiety: a system that detects arousal may detect exercise, caffeine, panic, startle, or trauma recall unless context is modeled explicitly.

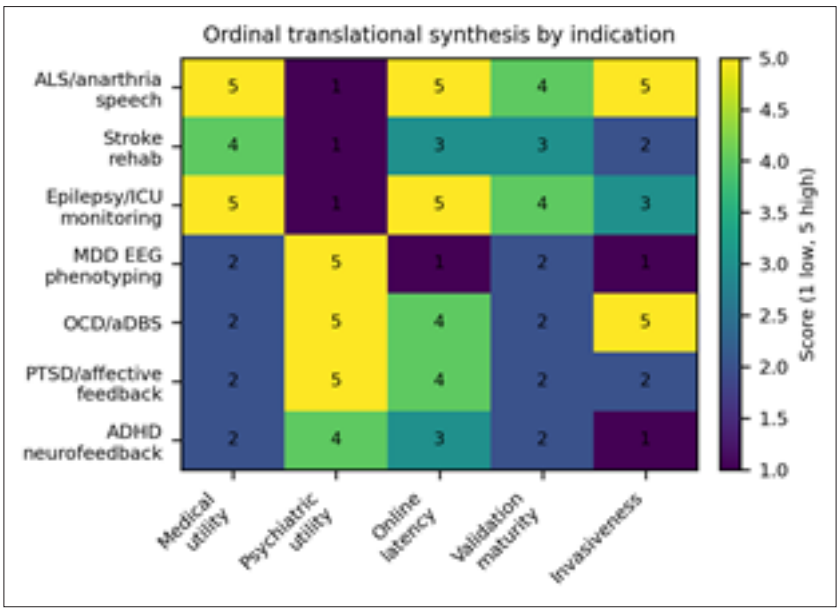

Figure 4: Ordinal translational synthesis for medical and psychiatric BCI indications. Scores are author-coded from the reviewed literature for visualization only; they are not a meta-analysis or regulatory determination.

Table 4: Psychiatric BCI Targets and Biomarker Risks

Indication	Candidate neural features	Critical validation risks
MDD / anhedonia	Frontal asymmetry, spectral power, functional connectivity, task reactivity, sleep and circadian markers.	Diagnostic heterogeneity, medication effects, suicidality risk, fairness, prospective relapse prediction.
OCD	Cortico-striato-thalamo-cortical activity, beta/theta/gamma markers, compulsive-state dynamics.	Causal biomarker validation, stimulation target specificity, symptom provocation ethics.
PTSD / anxiety	Amygdala-prefrontal coupling, autonomic coupling, threat-reactivity markers, affective neurofeedback.	Context confounding, privacy of trauma triggers, closed-loop exposure safety.
ADHD	Attention-state rhythms, event-related potentials, vigilance and impulsivity markers.	Developmental variability, placebo response, school/home ecological validation.
Psychosis / hallucination	Auditory-cortex, salience-network, and cognitive-control markers; feedback for symptom regulation.	Risk of delusion reinforcement, consent, fluctuating capacity, hallucination distress endpoints.
Addiction / craving	Cue-reactivity, prefrontal control, reward-network dynamics, relapse-risk physiology.	False alerts, stigma, behavioral confounds, intervention timing.

Multimodal Non-Invasive and Wearable Clinical Systems

Wearable BCI systems are attractive for psychiatry because symptoms are contextual, longitudinal, and often absent during laboratory visits. EEG alone is vulnerable to motion, facial EMG, ocular artifacts, and low spatial resolution; however, multimodal fusion with fNIRS, EOG, EMG, PPG, electrodermal activity, inertial sensing, speech, sleep, and smartphone context can improve ecological validity. A 2025 systematic review of multimodal EEG-based ML in clinical applications retained 16 studies and reported that 11 used multimodal data to improve accuracy (Zhao et al., 2025).

Fusion can occur at signal, feature, latent, or decision level. The main technical challenge is asynchronous sampling: EEG dynamics occur at millisecond scale, fNIRS hemodynamics at seconds, autonomic physiology at seconds to minutes, and psychiatric symptoms at minutes to weeks. A multimodal BCI should therefore represent temporal hierarchy explicitly:

$$z_t = f_{\text{EEG}}(X_t) \| f_{\text{fNIRS}}(H_{t-L:t}) \| f_{\text{auto}}(a_{t-M:t}) \| f_{\text{context}}(c_{t-Q:t}) \quad (10)$$

The concatenation operator is schematic; attention or state-space fusion may be preferable. In medical rehabilitation, multimodal signals can estimate effort, fatigue, compensatory movement, and therapeutic engagement. In psychiatry, the same fusion strategy can monitor arousal, sleep disruption, social withdrawal, medication adherence proxies, and affective instability, but only with stringent consent boundaries and patient control.

Table 5: Multimodal Fusion Levels for Clinical BCI

Fusion level	Mechanism	Medical/psychiatric interpretation
Signal-level	Synchronize and concatenate raw or filtered streams.	Rarely robust outside laboratory; sensitive to sampling mismatch and artifacts.
Feature-level	Extract spectral, connectivity, hemodynamic, autonomic, and behavioral features before model fusion.	Clinically interpretable; may discard nonlinear interactions.
Latent-level	Modality-specific encoders project into shared representation with attention or contrastive losses.	Powerful for missing sensors; requires careful leakage control.
Decision-level	Independent models vote or calibrate posterior probabilities.	Easier validation; may fail when modalities disagree pathologically.
Control-level	Fusion model directly modulates feedback, task difficulty, or stimulation.	Highest clinical risk; requires hard safety constraints and audit logs.

Evaluation: From Benchmark Scores to Clinical Evidence

Open benchmarking has become central to EEG-BCI reproducibility. MOABB, Braindecode, OpenNeuro, and NEMAR supply infrastructure for reproducible pipelines and dataset reuse (Chevallier et al., 2024; Aristimunha et al., 2026; Braindecode contributors, 2026; OpenNeuro contributors, 2026; NEMAR contributors, 2026). This discipline is especially important for foundation models because pretraining data can overlap with test data, normalization can leak information across sessions, and subject-dependent splits can be mistaken for subject-independent generalization.

Clinical BCI evaluation should report at least six quantities:

- split level,
- calibration budget,
- online latency,
- adaptation schedule,
- uncertainty and abstention behavior, and
- functional or symptom endpoint. For psychiatric models, the endpoint must be clinically anchored: depression severity, relapse, symptom burden, anxiety reactivity, compulsive time, sleep continuity, or functioning. Classification of diagnosis is rarely enough.

For imbalanced clinical labels, balanced accuracy, AUROC, AUPRC, calibration error, and decision-curve analysis are more informative than raw accuracy. A classifier used to trigger intervention must be calibrated:

$$\mathrm{ECE} = \sum_{b=1}^{B} \frac{|S_b|}{n} \left|\mathrm{acc}(S_b) - \mathrm{conf}(S_b)\right| \tag{11}$$

Expected calibration error (ECE) is especially relevant for neuropsychiatric monitoring, where overconfident false positives may cause unnecessary clinician alerts and false negatives may miss deterioration. For assistive BCIs, latency and throughput are decisive. For rehabilitation, longitudinal clinical endpoints such as Fugl-Meyer score, activities-of-daily-living function, fatigue, and adherence should be paired with decoder metrics (Elashmawi et al., 2024), (Williams et al., 2025).

Table 6: Minimum Reporting Requirements for Translational BCI-Ml

Item	Minimum technical report	Clinical rationale
Split definition	Subject-dependent, cross-session, cross-subject, cross-site, or prospective.	Different split levels answer different clinical questions.
Calibration budget	Minutes, trials, labels, clinician time, and target-task burden.	Determines usability in ALS, stroke, pediatric, and psychiatric cohorts.
Latency	Window length, inference time, update rate, output delay, stimulation delay.	Speech, motor control, and closed-loop DBS require real-time constraints.
Adaptation	Offline, online, unsupervised, supervised, or clinician-gated.	Uncontrolled adaptation can destabilize therapy or encode artifacts.
Safety	Uncertainty, abstention, safe stop, dose limits, audit logs.	Essential for stimulation, mobility, and psychiatric risk escalation.
Privacy	Consent, de-identification, federated learning, DP, model access.	Neural data can reveal health, identity, affect, and intended action.

Neural Data Privacy, Security, and Regulatory Risk

Neural recordings are sensitive for reasons that exceed ordinary biometric data. They can encode neurological disease, sleep, attention, affect, medication, identity, seizure risk, and intended action. Speech BCIs add a unique risk: intended but unspoken language may be private even when the patient has consented to assistive communication. Consumer EEG and psychiatric wearables add a second risk: clinically meaningful inferences may be made outside clinical oversight.

Privacy-preserving BCI methods include federated learning, secure aggregation, split learning, differential privacy, encrypted inference, and synthetic data. None is a free solution. Federated gradients can leak information under gradient-inversion attacks (Geiping et al., 2020), differential privacy can reduce utility, and synthetic neural data can reproduce

biases or identifiable dynamics. A formal privacy objective is often written as differential privacy:

$$\Pr[M(D) \in S] \leq e^{\varepsilon}\Pr[M(D') \in S] + \delta \quad (12)$$

For adjacent datasets D and D'. The parameters epsilon and delta must be interpreted clinically: a low privacy budget is not useful if it destroys seizure detection or speech decoding, whereas high utility is not acceptable if a model leaks psychiatric status or inner speech. Colorado expanded biological-data protection in 2024 to include neural data, illustrating that governance is becoming explicit rather than optional (Colorado General Assembly, 2024).

Implantable systems also create long-term obligations. The patient may depend on a proprietary device for communication or autonomy; therefore, software update policy, cybersecurity, battery replacement, explantation, interoperability, and vendor continuity are clinical risks. The emerging concept of durable or forever access emphasizes that a BCI is not a one-time implant but a long-lived medical infrastructure (Ugur et al., 2024).

Research Agenda for 2026-2028

First, EEG foundation models require prospective, leakage-free, clinically stratified benchmarks. The benchmark should fix calibration budgets and compare against Riemannian, EEGNet-like, and task-specific baselines, not only against other large models (Kuruppu et al., 2025; Dai et al., 2026; Yang et al., 2026), (Xiong et al., 2025; Wang et al., 2026; Portmann et al., 2026).

Second, psychiatric BCIs need dimensional labels and causal validation. Models should predict symptom trajectories, relapse, treatment response, and within-patient state transitions; they should not rely only on categorical DSM-style labels. Closed-loop stimulation should be tested with explicit safety constraints and clinician-gated escalation.

Third, speech and motor neuroprostheses need longitudinal home-relevant endpoints. A system that performs well in a laboratory session must be evaluated across weeks to months for calibration burden, user fatigue, caregiver workload, error correction, and psychosocial utility (Angrick et al., 2024; Littlejohn et al., 2025; Wairagkar et al., 2025; Card et al., 2024; Willett et al., 2023; Metzger et al., 2023; Willsey et al., 2025; Karpowicz et al., 2025).

Fourth, multimodal wearable BCI should use missing-modality training, uncertainty decomposition, and context-aware confound control. Psychiatric applications must distinguish arousal, exertion, panic, medication effects, and trauma-triggered states rather than treating all physiological activation as the same target.

Fifth, privacy, consent, and model governance should be engineered into the training objective and deployment workflow. Model cards for neural decoders should report data provenance, demographic composition, clinical exclusions, calibration assumptions, known failure modes, and restrictions on secondary inference.

Conclusion

The 2024-2026 BCI literature demonstrates a transition from offline decoding experiments to adaptive, closed-loop medical systems. Foundation models, hybrid architectures, sequence transducers, high-density implants, latent-dynamics alignment, and wearable multimodal sensors are expanding what can be decoded and acted upon. Yet clinical translation depends on factors that conventional machine-learning leaderboards often underreport: calibration burden, online latency, drift, prospective validation, patient-centered utility, privacy, and safety.

The medical and psychiatric frontier is therefore not a single architecture. It is a disciplined systems problem in which neural representations, control policies, clinical endpoints, regulatory constraints, and patient autonomy must be optimized jointly. BCIs for speech, movement, rehabilitation, depression, OCD, PTSD, ADHD, and other psychiatric indications will succeed only when their machine-learning methods are evaluated as components of longitudinal therapeutic loops.

References

1. Ding, Y., Ma, X., Zhang, P., Tang, Y., Zhao, Z., Chen, D., Wu, D., & Tang, Z. (2026). "Neural decoding for EEG-BCI: from conventional machine learning to deep learning models." Brain Hemorrhages, DOI: https://doi.org/10.1016/j.hest.2026.01.002
2. Saibene, A., Ghaemi, H., & Dagdevir, E. (2024). "Deep learning in motor imagery EEG signal decoding: A systematic review." Neurocomputing, 610, 128577. DOI: https://doi.org/10.1016/j.neucom.2024.128577
3. Hameed, I. Khan, D. M., Ahmed, S. M., Aftab, S. S., & Fazal, H. (2025). "Enhancing motor imagery EEG signal decoding through machine learning: A systematic review of recent progress." Computers in Biology and Medicine, 185, 109534. DOI: https://doi.org/10.1016/j.compbiomed.2024.109534
4. Elashmawi, W. H., Ayman, A., Antoun, M., Mohamed, H., Mohamed, S. E., Amr, H., Talaat, Y., & Ali, A. (2024). "A comprehensive review on brain-computer interface (BCI)-based machine and deep learning algorithms for stroke rehabilitation." *Applied Sciences, 14*(14), 6347. DOI: https://doi.org/10.3390/app14146347
5. Williams, C., Anik, F. I., Hasan, M. M., Rodriguez-Cardenas, J., Chowdhury, A., Tian, S., He, S., & Sakib, N. (2025). "Advancing brain-computer interface closed-loop systems for neurorehabilitation: systematic review of AI and machine learning innovations in biomedical engineering." *JMIR Biomedical Engineering*, 10, e72218. DOI: https://doi.org/10.2196/72218
6. Jiang, W. B., Zhao, L. M., & Lu, B. L. (2024). "Large Brain Model for learning generic representations with tremendous EEG data in BCI," International Conference on Learning Representations.

7. Jiang, W. B., Wang, Y., Lu, B. L. & Li, D. (2025). "NeuroLM: A universal multi-task foundation model for bridging the gap between language and EEG signals," International Conference on Learning Representations, 2025; arXiv:2409.00101. DOI: https://doi.org/10.48550/arXiv.2409.00101
8. Kuruppu, G., Wagh, N., Kremen, V., Pati, S., Worrell, G., & Varatharajah, Y. (2025). "EEG foundation models: A critical review of current progress and future directions." arXiv:2507.11783. DOI: https://doi.org/10.48550/arXiv.2507.11783
9. Dai, R., Shenhua, D., Baoliang, L., & Weilong, Z. (2026). "EEG foundation models for brain-computer interfaces: progress and future directions." *Journal of Shanghai Jiaotong University (Science), 31*(3), 783-799. DOI: https://doi.org/10.1007/s12204-026-2922-0
10. Yang, L., Sun, Q., Li, A., & Van Hulle, M. M. (2026). "Are EEG foundation models worth it? Comparative evaluation with traditional decoders in diverse BCI tasks," International Conference on Learning Representations. https://openreview.net/pdf?id=5Xwm8e6vbh
11. Xiong, W., Li, J., Li, J., & Zhu, K. (2025). "EEG-FM-Bench: A comprehensive benchmark for the systematic evaluation of EEG foundation models." arXiv:2508.17742. DOI: https://doi.org/10.48550/arXiv.2508.17742
12. Portmann, H., & Morishima, Y. (2026). "Systematic review of self-supervised foundation models for brain network representation using electroencephalography." arXiv:2602.03269. DOI: https://doi.org/10.48550/arXiv.2602.03269
13. Wang, J., Zhao, S., Zhou, Y., Kang, Y., Li, S., & Pan, G. (2026). "DeeperBrain: A neuro-grounded EEG foundation model towards universal BCI," arXiv:2601.06134. DOI: https://doi.org/10.48550/arXiv.2601.06134
14. Liang, S., Li, L., Zu, W., Feng, W., & Hang, W. (2024). "Adaptive deep feature representation learning for cross-subject EEG decoding," *BMC Bioinformatics, 25*(1), 393. DOI: https://doi.org/10.1186/s12859-024-06024-w?urlappend=%3Futm_source%3Dresearchgate.net%26utm_medium%3Darticle
15. Otarbay, Z., & Kyzyrkanov, A. (2025). "Transfer learning for subject-independent motor imagery EEG classification using convolutional relational networks," *Frontiers in Neuroscience*, 19, DOI: https://doi.org/10.3389/fnins.2025.1691929
16. Wu, D. (2025). "Revisiting Euclidean alignment for transfer learning in EEG-based brain-computer interfaces." *Journal of Neural Engineering, 22*(3), 031005. DOI: https://doi.org/10.1088/1741-2552/addd49
17. Liao, W., & Wang, W. (2024). "EEGEncoder: Advancing BCI with transformer-based motor imagery classification," arXiv:2404.14869. DOI: https://doi.org/10.48550/arXiv.2404.14869
18. Klein, T., Minakowski, P., & Sager, S. (2025). "Flexible Patched Brain Transformer model for EEG decoding," *Scientific Reports, 15*(1), 10935. DOI:https://doi.org/10.1038/s41598-025-86294-3?urlappend=%3Futm_source%3Dresearchgate.net%26utm_medium%3Darticle
19. Altaheri, H., Karray, F., & Karimi, A. H. (2025). "Temporal convolutional transformer for EEG based motor imagery decoding," *Scientific Reports, 15*(1), 32959. DOI:https://doi.org/10.1038/s41598-025-16219-7?urlappend=%3Futm_source%3Dresearchgate.net%26utm_medium%3Darticle
20. Vafaei, E., & Hosseini, M. (2025). "Transformers in EEG analysis: A review of architectures and applications in motor imagery, seizure, and emotion classification," *Sensors, 25*(5), 1293. DOI: https://doi.org/10.3390/s25051293
21. Angrick, M., Luo, S., Rabbani, Q., Candrea, D. N., Shah, S., Milsap, G. W., Anderson, W. S., Gordon, C. R., Rosenblatt, K. R., Clawson, L., Tippett, D. C., Maragakis, N., Tenore, F. V., Fifer, M. S., Hermansky, H., Ramsey, N. F., & Crone, N. E. (2024). "Online speech synthesis using a chronically implanted brain-computer interface in an individual with ALS," *Scientific Reports*, 14, 9617. https://www.nature.com/articles/s41598-024-60277-2
22. Littlejohn, K. T., Cho, C. J., Liu, J. R., Silva, A. B., Yu, B., Anderson, V. R., Kurtz-Miott, C. M., Brosler, S., Kashyap, A. P., Hallinan, I. P., Shah, A., Tu-Chan, A., Ganguly, K., Moses, D. A., Chang, E. F., & Anumanchipalli, G. K. (2025). "A streaming brain-to-voice neuroprosthesis to restore naturalistic communication." *Nature Neuroscience, 28*(4), 902-912. DOI: https://doi.org/10.1038/s41593-025-01905-6?urlappend=%3Futm_source%3Dresearchgate.net%26utm_medium%3Darticle
23. Wairagkar, M., Card, N. S., Singer-Clark, T., Hou, X., Iacobacci, C., Miller, L. M., Hochberg, L. R., Brandman, D. M., & Stavisky, S. D. (2025). "An instantaneous voice-synthesis neuroprosthesis" *Nature, 644*(8075), 145-152. DOI: https://doi.org/10.1038/s41586-025-09127-3
24. Card, N. S., Wairagkar, M., Iacobacci, C., Hou, X., Singer-Clark, T., Willett, F. R., Kunz, E. M., Fan, C., Vahdati Nia, M., Deo, D. R., Srinivasan, A., Choi, E. Y., Glasser, M. F., Hochberg, L. R., Henderson, J. M., Shahlaie, K., Stavisky, S. D., & Brandman, D. M. (2024). "An accurate and rapidly calibrating speech neuroprosthesis," *New England Journal of Medicine, 391*(7), 609-618. DOI: https://doi.org/10.1101/2023.12.26.23300110
25. Willett, F.R., Kunz, E.M., Fan, C., Avansino, D. T., Wilson, G. H., Choi, E. Y., Kamdar, F., Glasser, M. F., Hochberg, L. R., Druckmann, S., Shenoy, K. V., & Henderson, J. M. (2023). "A high-performance speech neuroprosthesis" *Nature*, 620, 1031-1036. DOI: https://doi.org/10.1038/s41586-023-06377-x
26. Metzger, S. L., Littlejohn, K. T., Silva, A. B., Moses, D. A., Seaton, M. P., Wang, R., Dougherty, M. E., Liu, J. R., Wu, P., Berger, M. A., Zhuravleva, I., Tu-Chan, A., Ganguly, K., Anumanchipalli, G. K., & Chang, E. F. (2023). "A high-performance neuroprosthesis for speech decoding and avatar control." Nature, 620, 1037-1046. DOI: https://doi.org/10.1038/s41586-023-06443-4
27. Willsey, M. S., Shah, N. P., Avansino, D. T., Hahn, N. V., Jamiolkowski, R. M., Kamdar, F. B., Hochberg, L. R., Willett, F. R., & Henderson, J. M. (2025). "A high-performance brain-computer interface for finger decoding and quadcopter game control in an individual with

paralysis." *Nature Medicine, 31*(1), 96-104.
DOI: https://doi.org/10.1038/s41591-024-03341-8
28. Karpowicz, B. M., Ali, Y. H., Wimalasena, L. N., Sedler, A. R., Keshtkaran, M. R., Bodkin, K., Ma, X., Rubin, D. B., Williams, Z. M., Cash, S. S., Hochberg, L. R., Miller, L. E., & Pandarinath, C. (2025). "Stabilizing brain-computer interfaces through alignment of latent dynamics." *Nature Communications, 16*(1), 4662.
DOI: https://doi.org/10.1038/s41467-025-59652-y
29. Nassibi, A., Papavassiliou, C., Rakhmatulin, I., Mandic, D., & Atashzar, S. F. (2025). "A systematic review of EEG-based machine intelligence algorithms for depression diagnosis, and monitoring," arXiv:2503.19820.
DOI: https://doi.org/10.48550/arXiv.2503.19820
30. Zhao, S., Li, W., Wang, X., Foglia, S., Tan, H., Zhang, B., Hamoodi, A., Nelson, A., & Gao, Z. (2025). "A systematic review of machine learning methods for multimodal EEG data in clinical application." arXiv:2501.08585.
DOI: https://doi.org/10.48550/arXiv.2501.08585
31. Chevallier, S., Carrara, I., Aristimunha, B., Guetschel, P., Sedlar, S., Lopes, S., Velut, S., Khazem, S., & Moreau, T. (2024). "The largest EEG-based BCI reproducibility study for open science: the MOABB benchmark," arXiv:2404.15319.
DOI: https://doi.org/10.48550/arXiv.2404.15319
32. Aristimunha, B. et al. (2026). "MOABB: Mother of all BCI Benchmarks, v1.5.0," Zenodo, 2026,
DOI: 10.5281/zenodo.10034223.
33. Braindecode contributors. (2026). "Braindecode: Decode raw EEG, ECoG and MEG with deep learning." open-source software and documentation.
DOI: https://doi.org/10.5281/zenodo.17699192
34. OpenNeuro contributors, "OpenNeuro: A free and open platform for sharing MRI, MEG, EEG, iEEG, and ECoG data," accessed May 2026.
35. NEMAR contributors. (2026). "NEMAR: NeuroElectroMagnetic data archive and tools for OpenNeuro," accessed May 2026.
36. Geiping, J., Bauermeister, H., Droge, H., & Moeller, M. (2020). "Inverting gradients: How easy is it to break privacy in federated learning?" Advances in Neural Information Processing Systems, 33. https://proceedings.neurips.cc/paper/2020/hash/c4ede56bbd98819ae6112b20ac6bf145-Abstract.html
37. Colorado General Assembly. (2024). "HB24-1058: Protect Privacy of Biological Data," approved Apr. 17, 2024; effective Aug. 7, 2024.
https://leg.colorado.gov/bills/hb24-1058
38. Ugur, M., Pothukuchi, R. P., & Bhattacharjee, A. (2024). " Towards Forever Access for Implanted a. Brain-Computer Interfaces," arXiv:2409.17496.
https://arxiv.org/pdf/2409.17496
39. Weng, W., Gu, Y., Guo, S., Ma, Y., Yang, Z., Liu, Y., Chen, Y. (2025). "Self-supervised learning for electroencephalogram: A systematic survey," arXiv:2401.05446, 2024; *ACM Computing Surveys, 57*(12), 1-38. DOI: https://doi.org/10.1145/3736574
40. Guetschel, P., Ahmadi, S., & Tangermann, M. (2024). "Review of deep representation learning techniques for brain-computer interfaces." *Journal of Neural Engineering, 21*(6), DOI: https://doi.org/10.1088/1741-2552/ad8962?urlappend=%3Futm_source%3Dresearchgate.net%26utm_medium%3Darticle
41. Yue, T., Gao, X., Xue, S., Tang, Y., Guo, L., Jiang, J., & Liu, J. (2024). "BrainGPT: Unleashing the potential of EEG generalist foundation model by autoregressive pre-training," arXiv:2410.19779.
DOI: https://doi.org/10.48550/arXiv.2410.19779
42. Ortega Caro, J., de O. Fonseca, A. H., Averill, C., Rizvi, S. A., Rosati, M., Cross, J. L., Mittal, P., Zappala, E., Levine, D., Dhodapkar, R. M., Abdallah, C. G., & Dijk, D. v. (2024). "BrainLM: A foundation model for brain activity recordings." International Conference on Learning Representations.
DOI: https://doi.org/10.1101/2023.09.12.557460
43. Yang, C., Westover, M., & Sun, J. (2023). "BIOT: Biosignal Transformer for cross-data learning in the wild," Advances in Neural Information Processing Systems.
DOI: https://doi.org/10.52202/075280-3420
44. Wang, S., Song, X., Song, X., Gu, Y., Cong, Z., Shen, Y., & Yu, L. (2026). "Non-invasive brain-computer interfaces: converging frontiers in neural signal decoding and flexible bioelectronics integration." *Nano-Micro Letters, 18*(1), 193. DOI: https://doi.org/10.1007/s40820-025-02042-2
45. Codina, T., Blankertz, B., & von Lühmann, A. (2025). "Multimodal fNIRS-EEG sensor fusion: Review of data-driven methods and perspective for naturalistic brain imaging," Imaging Neuroscience, 3,
DOI: https://doi.org/10.1162/imag.a.974
46. Gordon, E. C., & Seth, A. K. (2024). "Ethical considerations for the use of brain-computer interfaces for cognitive enhancement." PLOS Biology, 22(10), e3002899.
DOI: https://doi.org/10.1371/journal.pbio.3002899
47. Xia, K., Duch, W., Sun, Y., Xu, K., Fang, W., Luo, H., Zhang, Y., Sang, D., Xu, X., Wang, F. Y., & Wu, D. (2024). "Privacy-preserving brain-computer interfaces: A systematic review." arXiv:2412.11394.
DOI: https://doi.org/10.48550/arXiv.2412.11394
48. ClinicalTrials.gov. (2026). "PRIME: Precise Robotically Implanted Brain-Computer Interface." NCT06429735. https://clinicaltrials.gov/study/NCT06429735
49. ClinicalTrials.gov. (2026). "COMMAND Early Feasibility Study." NCT05035823. https://clinicaltrials.gov/study/NCT05035823
50. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). "Attention is all you need." Advances in Neural Information Processing Systems,
DOI: https://doi.org/10.48550/arXiv.1706.03762
51. Schirrmeister, R. T., Springenberg, J. T., Fiederer, L. D. J., Glasstetter, M., Eggensperger, K., Tangermann, M., Hutter, F., Burgard, W., & Ball, T. (2017). "Deep learning with convolutional neural networks for EEG decoding and visualization." Human Brain Mapping, 38(11), 5391-5420. DOI: https://doi.org/10.1002/hbm.23730

52. Lawhern, V. J., Solon, A. J., Waytowich, N. R., Gordon, S. M., Hung, C. P., & Lance, B. J. (2018). "EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces." *Journal of Neural Engineering, 15*(5), 056013. DOI: https://doi.org/10.1088/1741-2552/aace8c
53. Lotte, F., Bougrain, L., Cichocki, A., Clerc, M., Congedo, M., Rakotomamonjy, A., & Yger, F. (2018). "A review of classification algorithms for EEG-based brain-computer interfaces: a 10 year update." *Journal of Neural Engineering, 15*(3), 031005.
DOI: https://doi.org/10.1088/1741-2552/aab2f2
54. Dold, M., Pereira, J., Sajonz, B., Coenen, V. A., Thielen, J., Janssen, M. L. F., & Tangermann, M. (2024). "Dareplane: A modular open-source software platform for BCI research with application in closed-loop deep brain stimulation." *J Neural Eng, 22*(2).
DOI: https://doi.org/10.1088/1741-2552/adbb20
55. Sirbu, R., Morley, J., Schroder, T., Taddeo, M., Pothukuchi, R. P., Ugur, M., Bhattacharjee, A., & Floridi, L. (2025). "Regulating next-generation implantable brain-computer interfaces: Recommendations for ethical development and implementation," arXiv:2506.12540.
DOI: https://doi.org/10.48550/arXiv.2506.12540