Natural Resources for Human Health
eISSN: 2583-1194

**Original Research**

**Received** 28 March 2026
**Revised** 26 April 2026
**Accepted** 15 May 2026

Natural Resources for Human Health 2026, Vol. 6 (2s): 348–369

**KEYWORDS**:
healthcare, brain-computer interface, electroencephalography, neural decoding, deep learning, neurosecurity, differential privacy, federated learning,

# The Evolution and Security of Neural Communication Systems in Healthcare: From Electroencephalography to AI-Driven Brain-Computer-IoT Convergence

Shilpa Bhairanatti[1], Manjula Emmi[2], Rajeshwari Kisan[3], Chaitra Y Navalagi[4], Shridevi Hasanannavar[5]

[1,2]Associate Professor, S. G. Balekundri Institute of Technology, Belagavi-590010, Karnataka, India
[3,4,5] Assistant Professor, S. G. Balekundri Institute of Technology, Belagavi-590010, Karnataka, India

**ABSTRACT:** Neural communication systems, which carry information out of the nervous system without using nerves or muscles, have moved in a century and a half from a curiosity of animal electrophysiology to clinically deployed speech prostheses that operate unsupervised in the home, placing them at the centre of modern healthcare. This survey traces that trajectory across seventy primary sources and organises it into four eras: discovery and instrumentation, paradigm definition, clinical translation with statistical learning, and the present era of deep learning, foundation models and convergence with the Internet of Things. We compare acquisition modalities on a resolution against invasiveness trade space, chart the decoding literature from common spatial patterns and Riemannian geometry to large pretrained brain models, and quantify assistive throughput, which has risen from roughly two characters per minute in 1988 to a sustained fifty-six words per minute in long-term home use. We then argue that the same convergence responsible for these gains has multiplied the attack surface. Drawing on the classical side-channel literature and on brain-specific results, we propose a four-family taxonomy of leakage covering physical emanation, protocol metadata, semantic elicitation and learned-model channels, and map each onto the processing pipeline. Countermeasures are assessed as a five-tier defence-in-depth stack. We conclude that decoding capability now substantially outpaces the security evidence base, and set out a staged research agenda to close that gap before healthcare deployment scales further.

## 1. INTRODUCTION

A neural communication system is any engineered channel that conveys information from the nervous system to the outside world without passing through the ordinary output pathway of peripheral nerves and muscles. The definition is deliberately broad, covering the scalp electrode that resolves an intention to move a hand, the subdural grid that reconstructs an attempted sentence, and the endovascular array that lets a person with paralysis operate a computer. What unites them is a common structure: a transducer, a chain of signal conditioning and representation, a decoding stage that maps neural state onto symbols, and an actuator that renders those symbols in the world.

For most of its history the field advanced by improving one stage at a time: better electrodes, better spatial filters, better classifiers. Progress was real but incremental, and the field's own reviews were candid that information transfer rates of ten to twenty-five bits per minute placed practical communication far out of reach for the people who needed it most [1].

The last eight years have broken that pattern. Three developments arrived close enough together to compound. First, deep learning replaced hand-designed features with representations learned end to end from raw signals, then replaced task-specific training with self-supervised pretraining on unlabelled corpora spanning thousands of recordings [2-3]. Second, electrode technology reached channel counts that were implausible a decade ago, from ninety-six intracortical sites to arrays of a thousand or more surface contacts delivered without a craniotomy [4]. Third, the resulting systems stopped being laboratory instruments. They acquired radios, cloud backends, mobile applications and connections to the Internet of Things, and they left the clinic to enter mainstream healthcare [5-6].

The consequences for capability are well documented. The consequences for security are not. A system that is wireless, always on, connected to a remote decoder and physically attached to a person is an unusually rich target. It emits power and electromagnetic signatures that betray what its decoder is computing, transmits encrypted telemetry whose timing and volume betray what its user is doing, and runs a classifier that can be fooled by perturbations too small to see. Unlike a compromised password, the information it carries cannot be reissued.

The security literature has responded unevenly. The demonstration that a consumer headset could probe for a user's bank, PIN digits and recognised faces is more than a decade old [7], and the vocabulary of neurosecurity is older still [8]. Yet the field has no shared threat model, no standard adversarial benchmark, and no published measurement of power or electromagnetic leakage from a fielded brain-computer interface. The classical side-channel literature that would supply the method is mature [9-10]; it has simply not been pointed at this class of device.

This survey therefore has three aims: to organise the evolution of these systems into a periodisation that quantifies what each era delivered (Sections 2 and 3); to compare acquisition modalities, decoding families and applications on consistent axes (Sections 4 to 7); and to characterise the resulting security posture using a taxonomy specific to neural interfaces rather than borrowed from general computer security (Sections 8 and 9). Section 10 identifies the gaps and proposes a staged roadmap, and Section 11 concludes.

## 2. SURVEY SCOPE AND METHOD

Sources were gathered from indexed archival journals and conference proceedings, principally Nature and its family titles, Science, The New England Journal of Medicine, the relevant IEEE Transactions, the Journal of Neural Engineering, and the proceedings of the USENIX Security Symposium, the ACM Conference on Computer and Communications Security and the IEEE Symposium on Security and Privacy. Standards documents and regulatory guidance are cited where no peer-reviewed equivalent exists, and identified as such.

Three inclusion rules were applied. Historical sources were admitted when they are the primary record of a first demonstration. Technical sources were admitted when they report a quantified result on a named dataset or cohort. Security sources were admitted when they describe a demonstrated attack or a deployed defence rather than a hypothesised one. Every reference was verified against a publisher, PubMed or Crossref record; where a digital object identifier could not be confirmed it is omitted rather than inferred.

Two boundaries are worth stating. The survey covers communication out of the nervous system and the feedback that closes the loop, and treats therapeutic neuromodulation only where it creates a security surface. It excludes widely discussed systems with no peer-reviewed clinical publication: Neuralink has published its hardware platform [11] but no clinical results, and several reported implant milestones rest on press material. Communication rates are converted throughout to a words-per-minute equivalent at five characters per word, a coarse normalisation used for trend visualisation only.

## 3. Four Eras of Neural Communication

Figure 1 presents the periodisation used in this survey of healthcare-oriented neural communication. The boundaries are drawn where the field's dominant question changed rather than where a single paper appeared.

Natural Resources for Human Health
eISSN: 2583-1194

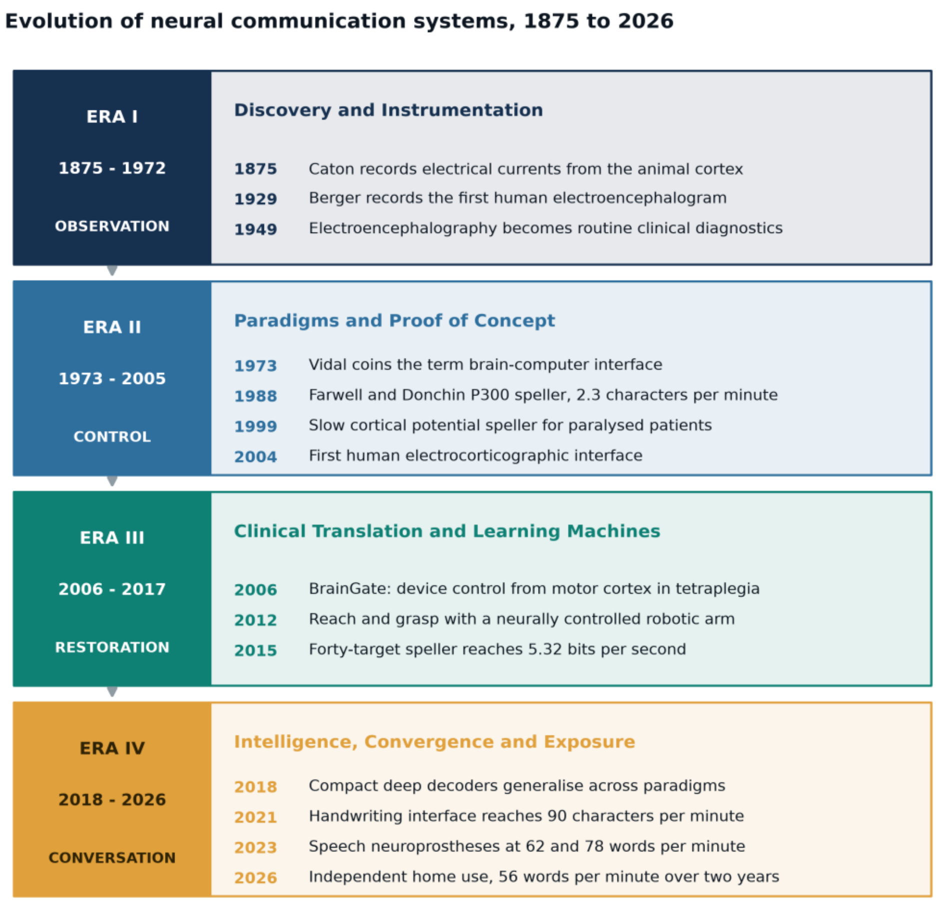


**Figure 1.** Evolution of neural communication systems, 1875 to 2026.

### 3.1 Era I: Discovery and Instrumentation, 1875 to 1972

The era opens with Caton's 1875 report of electrical currents recorded from the exposed hemispheres of rabbits and monkeys, and turns on Berger's 1929 recording of the human electroencephalogram, which named the signal and identified the alpha and beta rhythms [12]. The question of the period was whether cortical electrical activity could be recorded at all, reliably and repeatably; by mid-century the answer was settled well enough that electroencephalography had become routine clinical diagnostics. No system in this era attempted to decode. The capability was observation.

### 3.2 Era II: Paradigms and Proof of Concept, 1973 to 2005

Vidal's 1973 paper asked whether observable electroencephalographic signals could carry information in man-computer communication, coining the term brain-computer interface [13]. The following three decades produced the paradigm set that still structures the field. Farwell and Donchin built the P300 matrix speller, detecting the evoked response to an attended letter in a flashing character grid, at about 2.3 characters per minute [14]. Birbaumer and colleagues drove a binary spelling program for completely paralysed patients from self-regulated slow cortical potentials, at roughly two characters per minute [15]. Pfurtscheller and Lopes da Silva formalised event-related desynchronisation and synchronisation [16], the physiological substrate of every motor imagery interface, while steady-state visual evoked potential control emerged as the fourth durable paradigm.

The invasive branch matured in parallel, moving from single-electrode implants that recorded stable action potentials in locked-in patients to closed-loop ensemble control of robot arms in rodents and primates. Leuthardt demonstrated the first human electrocorticographic interface, with four epilepsy patients reaching seventy-four to one hundred percent success in a binary cursor task [17]. The era's self-understanding was consolidated by the 2002 review, which supplied the working definition of a brain-computer interface as a non-muscular channel for messages and commands and reported information transfer rates of ten to twenty-five bits per minute [1]. Shared platforms and the BCI Competition datasets then made cross-laboratory comparison possible and drew machine-learning researchers into the field.

### 3.3 Era III: Clinical Translation and Learning Machines, 2006 to 2017

The defining event is the BrainGate pilot trial, in which a ninety-six-microelectrode array in the motor cortex of a tetraplegic participant supported cursor control, television operation, prosthetic hand actuation and robot arm tasks [18]. Six years later two participants performed three-dimensional reach and grasp with robotic arms; one, using an array implanted more than five years earlier, drank from a bottle unaided for the first time in the fourteen years since her brainstem stroke [19]. The bottlenecks named at the time, implantable biocompatible hardware, real-time algorithms, sensory feedback and usable prostheses, would define the following decade.

On the algorithmic side this was the era of principled statistical learning. Common spatial patterns learned filters maximising variance for one motor imagery class while minimising it for the other [20], and filter-bank common spatial patterns added subject-specific frequency selection to win BCI Competition IV with a mean kappa of 0.569 [21]. Riemannian geometry then treated trial covariance matrices as points on a manifold, raising four-class accuracy on the same dataset from 65.1 to 70.2 percent [22]. The steady-state branch produced the era's throughput record: a forty-target speller using joint frequency-phase modulation reached 5.32 bits per second, about sixty characters per minute [23]. The capability of the era was restoration.

### 3.4 Era IV: Intelligence, Convergence and Exposure, 2018 to 2026

The present era is characterised by three simultaneous shifts, treated at length in Sections 5, 6 and 7: decoders became learned rather than designed, systems became networked rather than standalone, and communication rates crossed from assistive novelty into conversational usefulness. The capability is conversation, and the accompanying property, largely unaddressed, is exposure.

**Table 1.** The four eras of neural communication systems.

| Era | Period | Defining question | Enabling technology | Representative system | Limiting factor |
|---|---|---|---|---|---|
| I. Discovery | 1875-1972 | Can activity be recorded? | Valve amplifiers, ink recorders | Clinical electroencephalog raphy | Instrumentation noise |
| II. Paradigms | 1973-2005 | Can a signal carry a message? | Evoked-potential averaging, digital filtering | P300 speller, Graz motor imagery | Signal-to-noise, user training |
| III. Translation | 2006-2017 | Can it work in a patient? | Chronic arrays, spatial and manifold filtering | BrainGate, filter-bank spatial patterns | Stability, calibration burden |
| IV. Intelligence | 2018-2026 | Can it be fast and safe? | Deep learning, foundation models, edge computing | Speech neuroprostheses, AI-BCI-IoT platforms | Generalisation, governance, security |

Figure 2 renders the canonical signal chain that all four eras share, annotated with the trust boundaries that Section 8 will exploit.

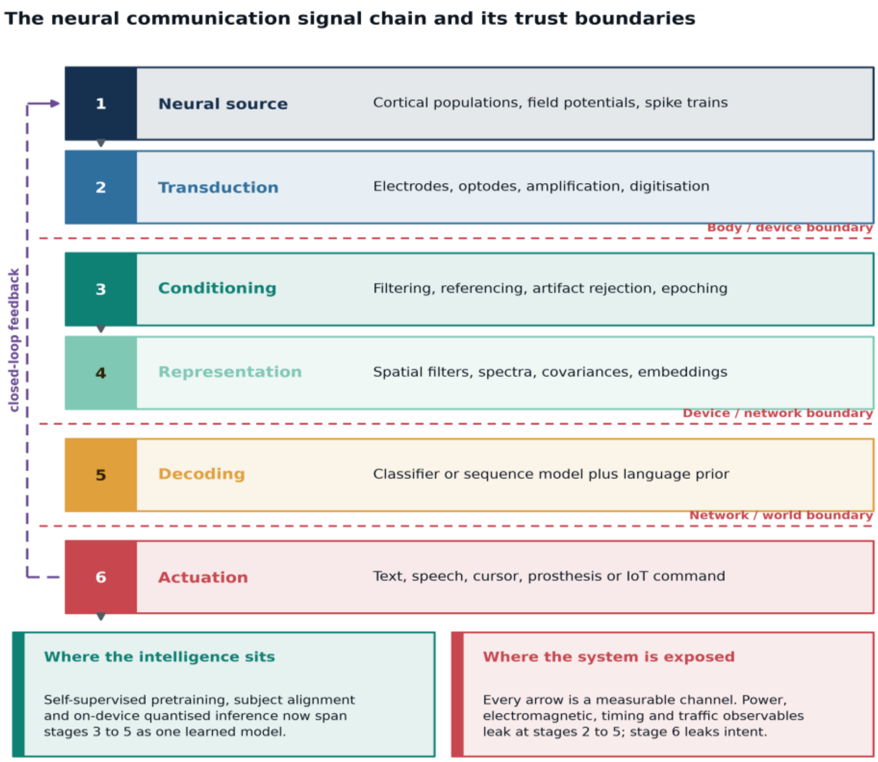


**Figure 2.** Canonical neural communication signal chain and its trust boundaries.

## 4. THE PHYSICAL LAYER: ACQUISITION MODALITIES AND INTERFACE HARDWARE

Every downstream healthcare claim rests on what the transducer can see. Figure 3 places the principal modalities on a trade space of effective spatial resolution against invasiveness, with marker area scaled by simultaneous channel count.

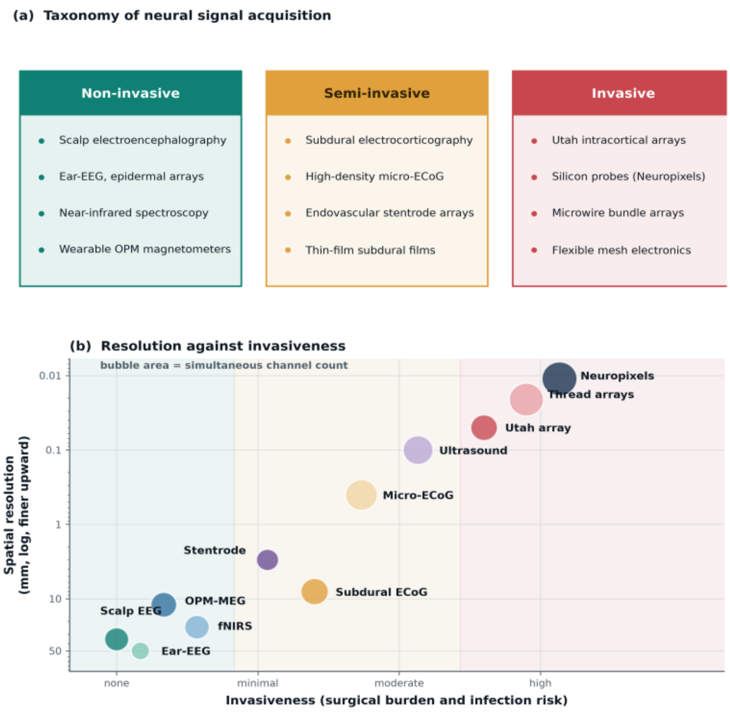


**Figure 3.** Acquisition modality taxonomy and the resolution-invasiveness trade space.

### 4.1 The invasive frontier

The Utah array remains the reference intracortical device, and its longevity is now quantified rather than assumed. An analysis of roughly six thousand recording datasets over nine years, spanning seventeen macaques and two humans, found a mean usable unit-recording lifespan of 622 days, with iridium oxide tips outperforming platinum [24]. Channel counts have since risen sharply. Neuropixels 2.0 places more than five thousand recording sites on a probe light enough that two probes plus headstage weigh just over one gram, with drift correction allowing the same neurons to be tracked for over two months, and the design has since been carried into acute human intraoperative recording [25]. Microwire bundles bonded to complementary metal-oxide-semiconductor amplifier arrays have pushed simultaneous channel counts into the tens of thousands, while a third strategy targets the mechanical mismatch that drives gliosis, using ultraflexible mesh electronics to achieve chronic single-unit stability.

### 4.2 The semi-invasive middle ground

Electrocorticography has always occupied a useful position between scalp recording and penetrating electrodes, and it has become substantially sharper. Flexible high-density grids now resolve submillimetre cortical organisation intraoperatively in humans, roughly an order of magnitude finer than the centimetre pitch of clinical grids, and thin-film surface arrays of 1024 channels have been delivered through a sub-millimetre cranial slit rather than a craniotomy [4]. The endovascular route avoids opening the skull entirely: the Stentrode showed no device-related serious adverse events, vessel occlusion or migration at twelve months in four participants with severe paralysis, with stable mean signal bandwidth of 233 hertz [26].

### 4.3 Non-invasive advances

Non-invasive recording has improved less in resolution than in wearability, which for communication systems matters more than it first appears. Semi-dry electrodes combine the signal quality of wet electrodes with the setup speed of dry ones, conducting-polymer tattoo electrodes provide metal-free skin contact, and ear-electroencephalography now produces results comparable to conventional scalp recording [27]. Optically pumped magnetometers removed the cryogenic constraint from magnetoencephalography, yielding a wearable helmet that records millisecond-resolution activity while participants move freely [28]. Wearable high-density functional near-infrared spectroscopy and functional ultrasound extend the same trend, trading temporal bandwidth for depth and portability.

A caution belongs here. A scoping review of 916 studies using consumer-grade equipment found Emotiv devices in 67.7 percent and NeuroSky in 24.6 percent, with brain-computer interfaces the largest application at 44.5 percent, yet only 6.2 percent were validation work [29]. The hardware most people will actually wear is the hardware whose signal quality is least characterised, and whose security is least examined.

**Table 2.** Comparison of acquisition modalities for neural communication.

| Modality | Invasiveness | Channels | Resolution | Bandwidth | Stability | Limitation |
|---|---|---|---|---|---|---|
| Scalp electroencephalography | None | 8-256 | 20-50 mm | 100 Hz | Session | Volume conduction |
| Ear and epidermal electrodes | None | 2-16 | 30-60 mm | 100 Hz | Wearable | Limited coverage |
| Functional near-infrared spectroscopy | None | 20-100 | 20-30 mm | 1 Hz | Good | Haemodynamic lag |
| Optically pumped magnetoencephalography | None | 20-128 | 10-20 mm | 500 Hz | Session | Cost, shielding |
| Endovascular stentrode | Minimal | 8-16 | 2-5 mm | 250 Hz | Multi-year | Few channels |
| Subdural electrocorticography | Moderate | 64-256 | 5-10 mm | 500 Hz | Multi-year | Craniotomy |
| High-density micro-electrocorticography | Moderate | 1024-2048 | 0.2-1 mm | 1 kHz | Unproven | Early evidence |
| Utah intracortical array | High | 96-256 | 0.05 mm | 30 kHz | 622 days | Gliosis, decline |
| Silicon probes, microwire bundles | High | 384-65,536 | 0.02 mm | 30 kHz | Preclinical | Data rate, heat |

Natural Resources for Human Health
eISSN: 2583-1194

## 5. FROM HAND-CRAFTED FEATURES TO FOUNDATION MODELS

### 5.1 The classical pipeline and its ceiling

Through the 2000s and early 2010s decoding meant designing a representation and fitting a shallow classifier to it: independent component analysis for artifact removal, wavelet sub-band decomposition for time-frequency features, and linear discriminant or support vector classifiers on top. The 2018 survey of the decade to 2017 is the revealing document. It concluded that adaptive classifiers repeatedly beat non-adaptive ones, that Riemannian methods had reached the state of the art, and, notably, that deep learning had not yet shown convincing improvement [30].

Two structural problems bounded the era. The first was cross-subject and cross-session distribution shift, which prevented reuse of training data and motivated a large transfer-learning literature. The second was the participation gap: an estimated fifteen to thirty percent of users could not reach the roughly seventy percent accuracy needed for practical control, a phenomenon partially addressed by co-adaptive schemes moving from a subject-independent to a subject-optimised classifier within one session [31]. Public benchmarks, notably the Temple University Hospital clinical corpus and the Tsinghua steady-state dataset, made both problems measurable and remain the substrate on which modern models are pretrained.

### 5.2 End-to-end learning

Two 2017 and 2018 papers reset the default. Deep and shallow convolutional networks decoded raw electroencephalography end to end, reaching 84.0 percent mean accuracy against 82.1 percent for filter-bank common spatial patterns, with input-perturbation visualisation confirming the networks had learned alpha, beta and high-gamma modulations consistent with known physiology [32]. EEGNet then showed that a compact depthwise and separable convolutional backbone generalises across four paradigms without architectural change [33], and it remains the standard lightweight baseline. Systematic reviews of the period are usefully unromantic: across 154 papers the median gain of deep learning over baselines was about five percent, and most released neither data nor code [2]. The MOABB framework responded directly, re-evaluating six pipelines across twelve open datasets and more than 250 subjects, and concluding that algorithms validated on a single dataset are not representative [34]. Architectural work since has followed the wider machine-learning field, through convolutional-transformer hybrids, graph networks over the electrode montage, diffusion-based augmentation under data scarcity, and compact temporal convolutions built for embedded deployment.

### 5.3 Self-supervision and brain foundation models

The decisive change is that decoders no longer need training from scratch on labelled data. Self-supervised pretext tasks applied to large clinical corpora produce representations on which linear classifiers beat fully supervised networks in low-label regimes. BENDR adapted the wav2vec 2.0 recipe to electroencephalography, pretraining a convolutional encoder and transformer contrastively on roughly 1.5 terabytes of clinical recordings from more than ten thousand people [35]. LaBraM pretrained on approximately 2500 hours drawn from around twenty heterogeneous datasets, solving the mismatched-montage problem through channel-patch tokenisation and reporting state-of-the-art results across abnormality detection, event classification, emotion recognition and gait prediction [3]. Successor models replace joint spatiotemporal attention with parallel criss-cross branches and montage-agnostic positional encoding.

Generative decoding has advanced in parallel and is directly relevant to Section 8. Latent diffusion models reconstruct high-resolution natural images from functional magnetic resonance imaging without fine-tuning the generative model itself. A language-model-based decoder trained on sixteen hours of narrative-listening data per participant recovers intelligible continuous word sequences for perceived speech, imagined speech and silent video [36]. That study also reports the most important privacy result in the modern decoding literature: decoding fails without the subject's cooperation during both training and application. Contrastive alignment of magnetoencephalography with self-supervised speech representations, across 175 volunteers, achieved up to 41 percent top-ten segment identification on average [37]. Explainability remains underdeveloped, with existing work targeting developers justifying predictions rather than clinicians or end users.

**Table 3.** Representative decoder families and reported performance.

| Family | Method | Benchmark | Reported result | Year |
|---|---|---|---|---|
| Spatial filtering | Common spatial | Two-class motor | 90.8-99.7%, three | 2000 |

| Family | Method | Benchmark | Reported result | Year |
|---|---|---|---|---|
| | patterns | imagery | subjects | |
| Filter-bank filtering | Filter-bank spatial patterns | Competition IV-2a | Mean kappa 0.569 | 2012 |
| Manifold learning | Tangent-space discriminant | Competition IV-2a | 70.2% four-class | 2012 |
| Correlation-based | Task-related components | Forty-target steady state | 89.8%, 325 bits/min | 2018 |
| End-to-end convolution | Deep ConvNet | Cross-dataset motor imagery | 84.0% versus 82.1% | 2017 |
| Compact convolution | EEGNet | Four paradigms | Generalises without redesign | 2018 |
| Embedded convolution | EEG-TCNet | Competition IV-2a | 77.4%, 83.8% tuned | 2020 |
| Convolution-transformer | EEG Conformer | Multiple public datasets | State of the art, interpretable | 2023 |
| Self-supervised | BENDR | Five downstream datasets | Competitive with supervised | 2021 |
| Foundation model | LaBraM | 2500 h pretraining | Best on four task families | 2024 |
| Foundation model | CBraMod | Ten tasks, twelve datasets | Best, montage-agnostic | 2025 |

## 6. HEALTHCARE APPLICATIONS

### 6.1 Assistive communication: the performance discontinuity

Figure 4 charts communication throughput and makes the discontinuity visible. Only studies that themselves report a rate are plotted; work reporting error rates or information transfer rates alone is discussed in the text but not converted, since such conversions are not comparable across paradigms.

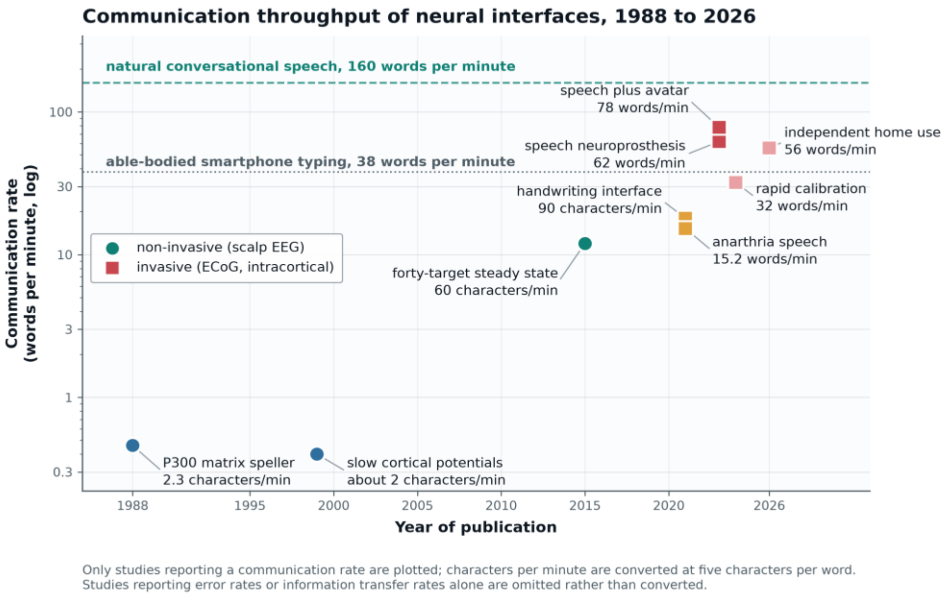


**Figure 4.** Communication throughput of neural interfaces, 1988 to 2026.

The invasive branch moved first. Anumanchipalli and colleagues established brain-to-voice synthesis, mapping cortical activity to vocal-tract kinematics and then to acoustics [38]. Moses reported the first clinical full-word speech decoding in a person with anarthria, at a median 15.2 words per minute with 25.6 percent word error on a fifty-word vocabulary and signals stable across 81 weeks [39]. Willett's handwriting interface took a different route, decoding attempted handwriting with a recurrent network at ninety characters per minute and 94.1 percent raw online accuracy [40]. Fixed vocabularies fell next, as silently attempted code-word spelling was shown to generalise beyond nine thousand words.

The 2023 results changed the scale of the problem. Willett's speech neuroprosthesis decoded attempted speech at 62 words per minute, 3.4 times the previous record, with 9.1 percent word error on fifty words and 23.8 percent on 125,000 [41]. Metzger's multimodal system produced text, synthesised speech matching the participant's pre-injury voice, and a facial avatar from a single decoder, at a median 78 words per minute [42]. Card then addressed calibration, reaching 99.6 percent on the first day of use after thirty minutes of calibration and sustaining 97.5 percent over 8.4 months [43]. Latency and expressiveness followed: streaming decoders emit speech in eighty-millisecond increments, and an instantaneous synthesiser decodes voice within ten milliseconds of the neural measurement, with listeners identifying the output at 100 percent median accuracy and the participant able to modulate pitch well enough to sing [44].

The most consequential recent datapoint is not a speed record. A participant with amyotrophic lateral sclerosis used an intracortical speech and cursor interface independently at home for more than 3800 hours across nearly two years, producing 183,060 sentences at an average 56 words per minute, 92 percent rated at least mostly correct [6]. This is the transition from demonstration to appliance, and precisely the transition that makes Section 8 urgent.

Non-invasive systems have not kept pace, and honesty about the gap matters. Brain2Qwerty decodes memorised typed sentences from magnetoencephalography at 32 percent character error but from electroencephalography at 67 percent [45]. Covert speech decoding from scalp recordings remains immature for real-time use, and for completely locked-in patients the achievable rate is around one character per minute even with intracortical electrodes.

### 6.2 Sign language and gesture-based inclusive communication

A parallel literature addresses communication for people with hearing and speech impairment, sharing the decoding stack while avoiding the neural interface entirely. Vision-based recognition was reset by large gloss-level benchmarks and then by sign language transformers, which learn recognition and translation jointly end to end and more than doubled translation quality on RWTH-PHOENIX-Weather-2014T from 9.58 to 21.80 BLEU-4 [46]. Sensor-based approaches trade generality for accuracy, with stretchable strain-sensor gloves converting several hundred sign patterns to speech at above 98 percent in under a second, and silent-speech interfaces occupy a third position, reading facial neuromuscular activity during subvocal verbalisation. All three share the security properties analysed in Section 8 without carrying the mental-privacy exposure of a direct neural recording.

**Table 4.** Milestones in assistive neural and gestural communication.

| Year | System | Modality | Vocabulary | Rate or accuracy |
|---|---|---|---|---|
| 1988 | P300 matrix speller | Scalp electrodes | Character grid | 2.3 characters/min |
| 1999 | Slow cortical potential speller | Scalp electrodes | Binary support | About 2 characters/min |
| 2015 | Frequency-phase speller | Scalp electrodes | Forty targets | 5.32 bits/s |
| 2018 | AlterEgo silent speech | Facial neuromuscular | Task-limited | 92% word accuracy |
| 2020 | Sign-to-speech glove | Strain sensors | 660 gestures | 98.63%, under 1 s |
| 2021 | Handwriting neuroprosthesis | Intracortical | Alphabetic | 90 characters/min |
| 2021 | Anarthria speech prosthesis | Electrocorticography | 50 words | 15.2 words/min |
| 2022 | Completely locked-in speller | Intracortical | 30 items | 1.08 characters/min |
| 2023 | Speech neuroprosthesis | Intracortical | 125,000 words | 62 words/min |
| 2023 | Speech plus avatar | Electrocorticography | 1024 words | 78 words/min |
| 2024 | Rapidly calibrating speech | Intracortical | 125,000 words | 97.5%, 8.4 months |
| 2025 | Instantaneous voice synthesis | Intracortical | Open | 10 ms latency |
| 2026 | Independent home use | Intracortical | 125,000 words | 56 words/min, 3800 h |

eISSN: 2583-1194

### 6.3 Defence and cognitive enhancement

Military interest has shaped the field since its early years, and the resulting programmes are now explicit. The DARPA Next-Generation Nonsurgical Neurotechnology programme targeted bidirectional non-surgical interfaces reading from and writing to sixteen independent channels within sixteen cubic millimetres of tissue at fifty millisecond latency [47]. A RAND assessment of United States military applications is notable for framing interface compromise as an attack pathway into an operator's motor and emotional circuits, capable of false motor commands or induced confusion [48].

Demonstrated capability is more modest than the rhetoric. Continuous three-dimensional quadcopter control by sensorimotor-rhythm modulation was shown in five trained subjects acquiring up to 90.5 percent of targets [49], and later steady-state systems extended this to swarm formation control at around ninety percent accuracy. Passive monitoring is closer to deployment, with small wireless headsets classifying pilot mental workload in the high eighties across thousands of epochs, within the broader neuroergonomics agenda. Claims for cognitive enhancement should be read carefully. A systematic review of 34 randomised transcranial direct current stimulation trials found 28 reporting cognitive gains but concluded the technology is not yet recommendable for operational use [50], and a later update found some promise for multitasking but no dependable vigilance effects, most studies carrying low to moderate risk of bias.

## 7. CONVERGENCE WITH ARTIFICIAL INTELLIGENCE, THE INTERNET OF THINGS AND THE NETWORK EDGE

The systems of Era IV are not standalone instruments but nodes in a connected healthcare ecosystem. Figure 5 sets out a layered reference architecture that captures how a modern deployment is actually partitioned.

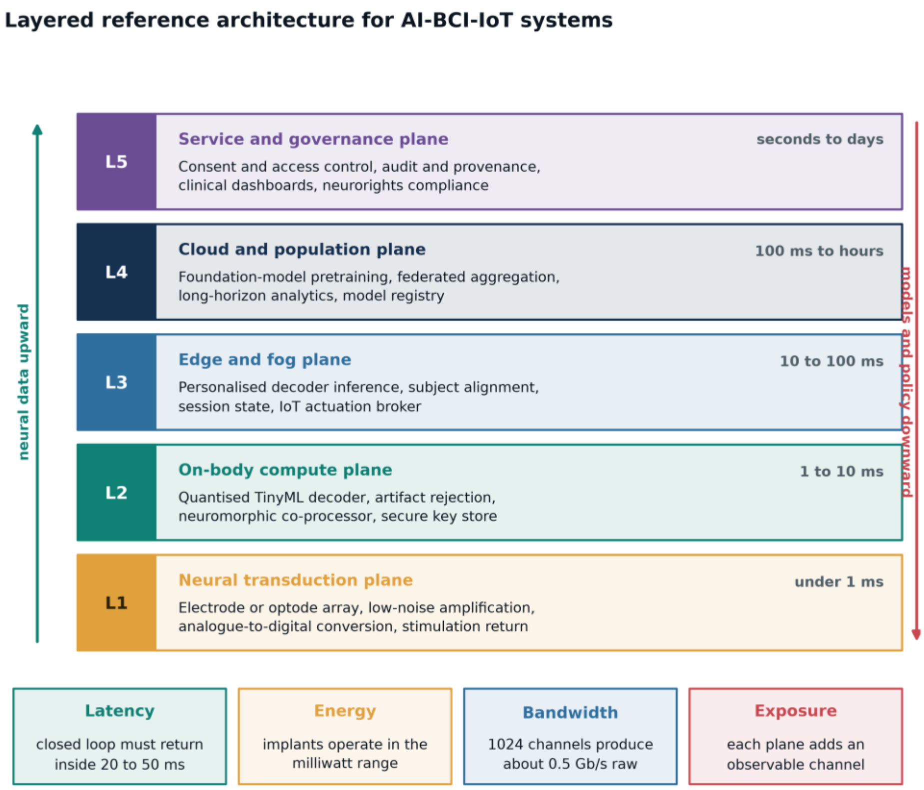


**Figure 5.** Layered reference architecture for convergent AI-BCI-IoT systems.

The application driver is direct actuation of connected objects. A unified deep-learning framework for human-thing cognitive interactivity established the brain-Internet-of-Things pattern, pairing a reinforcement-learned focal-zone selector with a recurrent classifier [5], and steady-state smart-home systems built on consumer headsets now actuate household appliances at close to ninety percent online accuracy. At the network layer, wireless body area networks supply the on-body link design, while the ultra-reliable low-latency budgets that closed-loop neural teleoperation imposes on sixth-generation networks are stated explicitly in the tactile-internet literature.

Where inference runs is the architecturally decisive question. Placing decoding on fog nodes rather than in the cloud cuts round-trip latency for real-time control. On-body inference has become genuinely practical: a quantised motor-imagery network with automatic channel selection achieves 82.5 percent two-class accuracy at thirty microjoules and 2.95 milliseconds per inference on an ultra-low-power microcontroller [51]. Neuromorphic co-integration goes further, combining recording and spiking-network inference on a single die at sub-milliwatt power, and at the opposite extreme personalised whole-brain simulation has been proposed as a substrate for brain-machine interfacing.

The security significance of this architecture is structural rather than incidental. Each additional plane introduces a physical observable, a network observable and an administrative domain. The energy and bandwidth constraints that force feature extraction on-body are the same constraints that make constant-time, masked or padded implementations expensive. Efficiency and exposure are coupled.

## 8. SECURITY AND PRIVACY: THE SIDE-CHANNEL THREAT

### 8.1 Why the general side-channel literature applies

The foundational results are unambiguous that correct implementations leak through physical observables. Differential power analysis extracts secrets from tamper-resistant devices even when the per-trace signal lies well below the noise floor [9], and the same principle has been carried through timing, electromagnetic, acoustic and microarchitectural channels, several of which recover complete key material on hardware that already carries countermeasures against power analysis. Encrypted traffic remains highly informative independently of any of this: convolutional models identify the site a user is visiting at over 98 percent accuracy from packet traces alone [52].

Three results move this directly onto the neural-interface substrate now deployed in healthcare. Power traces from an accelerator running a convolutional network recover the private input image at up to 89 percent accuracy without knowledge of the network parameters [53]. Passive timing and electromagnetic measurement of an ARM Cortex-M microcontroller, exactly the class of processor used in wearable decoders, recovers activation functions, layer and neuron counts and weights [10]. The machine-learning privacy literature supplies the rest: shadow-model membership inference determines whether a record was in a model's training set from black-box access alone, demonstrated against commercial services on a hospital discharge dataset, and model inversion recovers recognisable training examples from returned confidence values [54].

### 8.2 Neural-interface-specific threats

The brain-specific literature begins with neurosecurity, the application of computer-security principles to neural engineering, and with the argument that neither engineering practice nor medical trials nor neuroethical review ensures robustness against an adversary who alters, blocks or eavesdrops on neural signals [8]. That argument was concrete from the start, since wireless attacks on commercial implantable cardiac devices had already disclosed patient data and induced a command shock using an oscilloscope and a software radio.

The defining demonstration is Martinovic's side-channel attack, in which a malicious application with legitimate access to a consumer headset used P300 oddball stimuli to probe twenty-eight participants for bank cards, PIN digits, area of residence and known faces, reducing the entropy of that private information by approximately fifteen to forty percent relative to random guessing [7]. The threat model was then generalised to the application-store setting, with third-party applications extracting information users never intended to disclose and manufacturers prioritising functionality over security. For stimulating implants the term is brainjacking, distinguishing blind attacks such as halting stimulation or draining the battery from targeted attacks that impair motor function, alter impulse control or modulate the reward system [55].

Adversarial machine learning is now a mature sub-literature here. An unsupervised fast gradient sign method fools three widely used convolutional classifiers with small perturbations, and the attacks transfer without knowledge of the target's architecture, parameters or training data [56]. Perturbations too small to notice can drive P300 and steady-state spellers to output arbitrary attacker-chosen text, with consequences ranging from user frustration to severe misdiagnosis [57]. Backdoor attacks are worse operationally: a narrow period pulse injected as a poisoning trigger forces any test sample carrying the key into a target class, and need not be synchronised with the trial, which makes physical delivery far easier [58]. Finally, neural data is inherently re-identifiable. The CEREBRE protocol achieved 100 percent identification accuracy in a fifty-user pool from event-related potentials elicited across several functional systems [59], and resting-state recordings alone support inference of personality traits and executive function.

### 8.3 A taxonomy of leakage

Figures 6 and 7 organise this material. Figure 6 maps attack families onto the five pipeline stages; Figure 7 decomposes leakage into four families with their observables and disclosures.

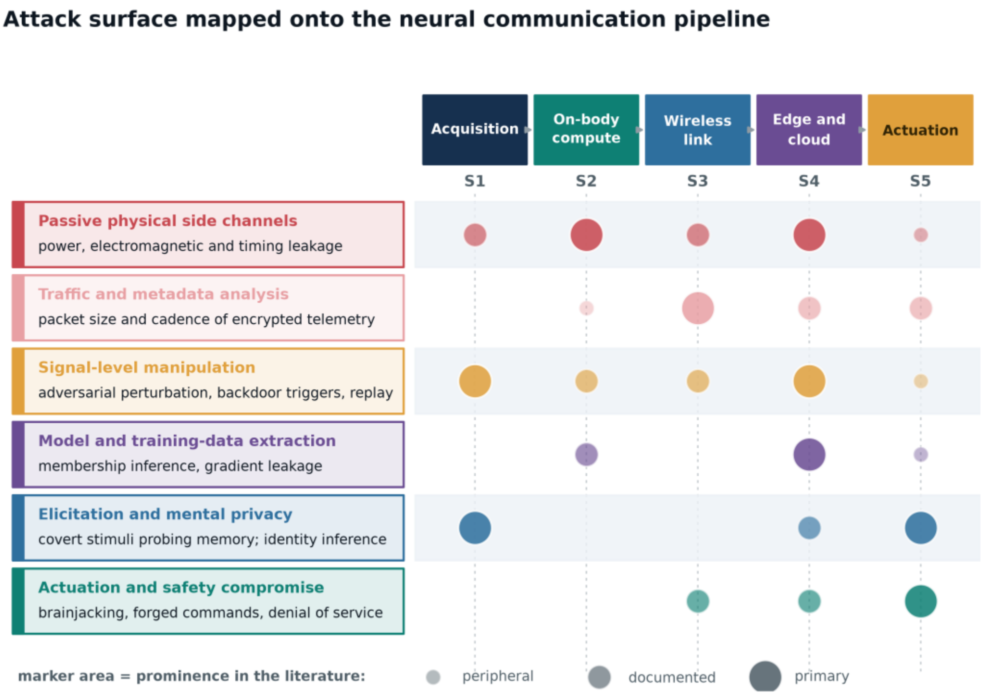


**Figure 6.** Attack surface mapped onto the neural communication pipeline.

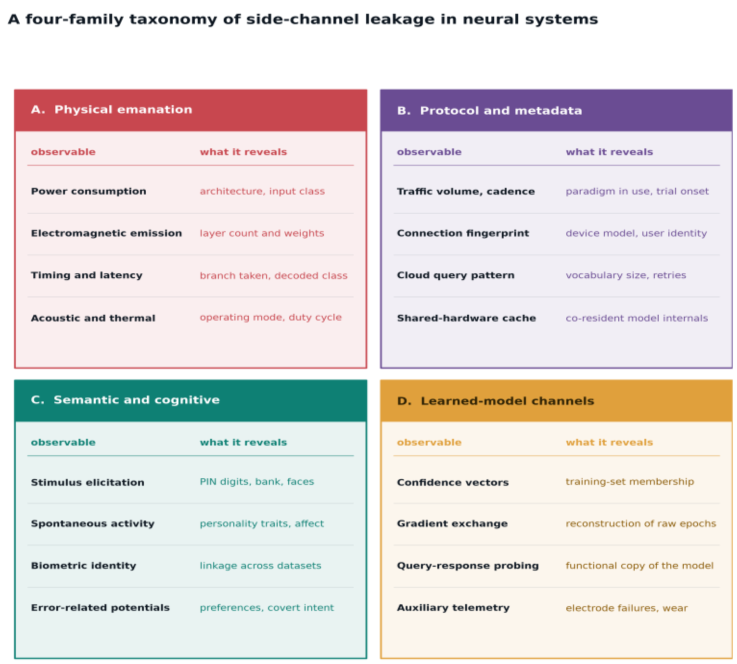


**Figure 7**. Four-family taxonomy of side-channel leakage in neural systems.

The taxonomy separates channels by what the adversary measures rather than by what they intend. Physical emanation channels are the classical family transplanted: supply current, near-field radiation, inference timing and package thermal behaviour, each shown to disclose model structure or private inputs on comparable hardware. Protocol and metadata channels exploit the fact that encryption protects content but not shape; packet cadence in an always-on telemetry stream marks trial onsets and session boundaries, and shared-tenancy edge nodes reintroduce cache contention. Semantic and cognitive channels are unique to this domain and the most consequential: the stimulus itself is the probe, and the leak is not a key but a memory. Learned-model channels attack the decoder as an asset, through confidence vectors, shared gradients, repeated queries and auxiliary telemetry such as impedance logs that disclose clinical state.

The asymmetry worth naming is that the first, second and fourth families have well-developed measurement methodologies that have simply not been applied to fielded neural hardware, while the third is where the harm is irreversible and the least methodological work exists.

Table 5. Attack taxonomy with representative demonstrated evidence.

| Family | Attack | Target stage | Demonstrated on | Reported effect |
|---|---|---|---|---|
| Physical emanation | Differential power analysis | Compute, radio | Tamper-resistant devices | Key recovery below noise |
| Physical emanation | Electromagnetic extraction | On-body compute | ARM Cortex-M | Layers, weights recovered |
| Physical emanation | Power analysis of inference | Edge accelerator | FPGA accelerator | Input recovered, up to 89% |
| Protocol metadata | Traffic fingerprinting | Wireless link | Tor traces | Over 98% identification |
| Protocol metadata | Cache contention | Shared edge node | Last-level cache | 96.7% of key bits |
| Signal manipulation | Adversarial perturbation | Acquisition, decoder | Three EEG classifiers | Transferable misclassification |
| Signal manipulation | Speller perturbation | Decoder | P300, steady-state spellers | Attacker-chosen output |
| Signal manipulation | Period-pulse backdoor | Training pipeline | Multiple EEG datasets | Target-class forcing |
| Model extraction | Membership inference | Cloud decoder | Clinical records | Training membership |
| Model extraction | Gradient inversion | Federated training | Image, text models | Pixel-accurate reconstruction |
| Semantic elicitation | P300 oddball probing | Acquisition, stimulus | 28 headset users | 15-40% entropy reduction |
| Semantic elicitation | Re-identification | Stored data | 50-user pool | 100% identification |
| Actuation compromise | Brainjacking | Stimulation path | Deep brain stimulators | Motor, affective modulation |

## 9. COUNTERMEASURES, STANDARDS AND GOVERNANCE

Figure 8 organises defences as a five-tier stack and, in panel (b), quantifies the cost that the strongest data-protection tier imposes.

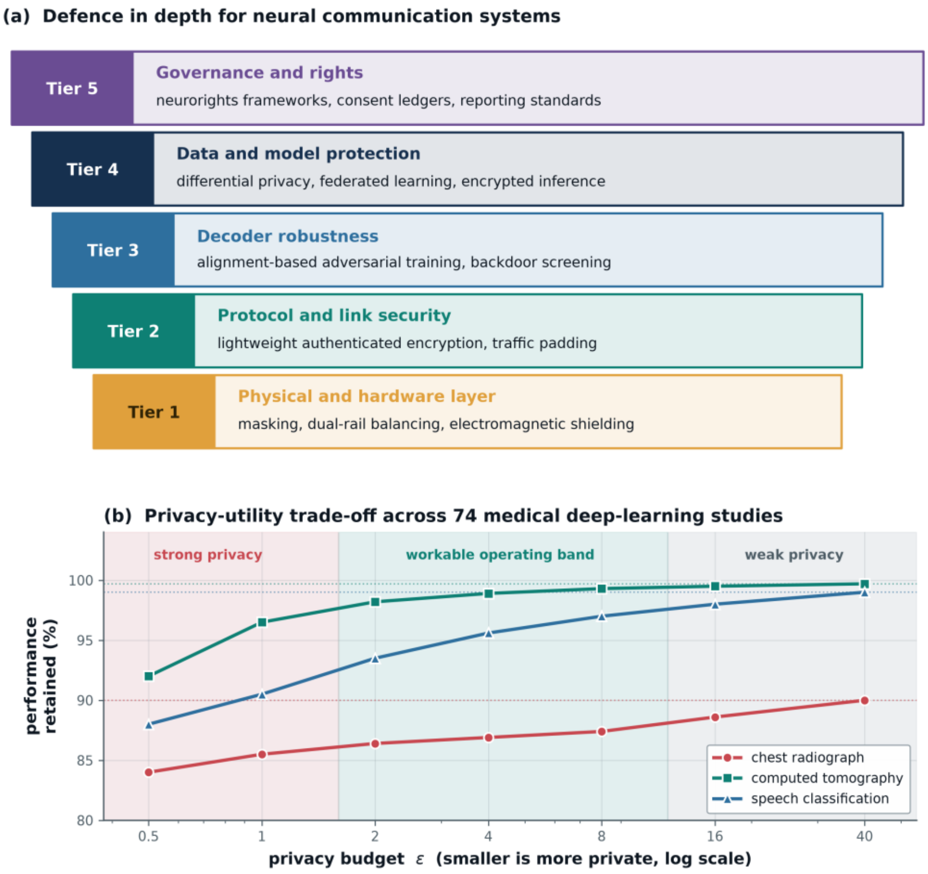


**Figure 8.** Defence-in-depth stack and the privacy-utility trade-off.

### 9.1 Hardware and link layers

The countermeasures for the physical family are classical and well founded. Masking transforms a circuit into a secret-shared equivalent provably secure against any fixed number of probes [60], and hiding approaches such as dual-rail precharge logic equalise per-cycle power consumption within standard design flows. Both cost area and energy, which is exactly what an implant does not have, and this constraint is the reason lightweight rather than conventional cryptography is required at the link layer; standardised compact authenticated-encryption primitives now exist for that role. Key establishment can exploit the biosignal itself, since electroencephalography is repeatable enough to regenerate a cryptographic key. Trusted execution environments warrant caution rather than reliance, since software-only transient-execution attacks have extracted enclave secrets and attestation keys.

### 9.2 Data, model and decoder protection

Differential privacy supplies the formal guarantee, and DP-SGD with per-example gradient clipping and the moments accountant made it practical for deep networks [61]. The cost is now quantified: a synthesis of 74 medical deep-learning studies found that a privacy budget of about ten preserves clinically acceptable accuracy while a budget of about one causes substantial loss, with chest radiograph area under the curve falling from 90 to 84 percent at a budget of 0.5 [62]. No comparable curve exists for neural data, which is the gap panel (b) of Figure 8 exposes. Federated learning keeps raw recordings local and works in medicine, with multi-hospital training matching pooled-data accuracy. Applied here, FedBS, using local batch-specific normalisation with sharpness-aware minimisation, outperforms six federated baselines and also exceeds centralised training that offers no privacy protection at all [63]. Federation alone is not a guarantee: shared gradients permit pixel-accurate reconstruction of training data [64], so secure aggregation, in which the server learns only the sum of client updates, is a necessary companion rather than an optional extra. Encrypted inference is viable for biosignals, and applied to electroencephalography for seizure detection it keeps inference in the millisecond range with plaintext-comparable performance, at a substantial training cost [65]. For decoder robustness, generic adversarial training degrades accuracy, whereas data alignment before adversarial training improves both robustness and clean accuracy across five datasets [66], and the first systematic defence benchmark concludes that combining augmentation, alignment and robust training beats any subset [67]. A taxonomy mapping attacks and defences across the full interface cycle organises this tier [68].

### 9.3 Standards, regulation and rights

The normative layer has moved faster than the engineering one. The neurorights framework proposes rights to cognitive liberty, mental privacy, mental integrity and psychological continuity [69], and the parallel statement of four ethical priorities, covering privacy and consent, agency and identity, augmentation and bias, set the agenda for the technical community. The warning that consumer-grade neurotechnology without safeguards enables subtle surveillance and manipulation remains the most precise statement of the gap. Formal standardisation is under way, with a functional model unifying interface terminology and minimum reporting requirements for in-vivo neural interface research [70], though neither yet requires disclosure of security posture. Regulatory instruments are converging from the general direction: cybersecurity guidance for healthcare medical devices now requires a secure product development framework, threat modelling and a software bill of materials in premarket submissions, and the European Union's artificial intelligence regulation prohibits emotion-inference systems in workplaces and education outside medical and safety purposes, a provision that bears directly on affective decoding.

**Table 6.** Countermeasures, their target, cost and current maturity.

| Tier | Countermeasure | Targets | Cost | Maturity here |
|---|---|---|---|---|
| Physical | Masking, secret sharing | Power, electromagnetic analysis | Area and energy | Untested on interfaces |
| Physical | Dual-rail power balancing | Differential power analysis | Double cell area | Untested on interfaces |
| Link | Lightweight authenticated encryption | Eavesdropping, injection | Small | Standard available |
| Link | Padding, constant cadence | Metadata analysis | Continuous energy | Not adopted |
| Decoder | Alignment-based adversarial training | Evasion | Modest, aids accuracy | Benchmarked |
| Decoder | Backdoor and trigger screening | Poisoning | Training time | Early |
| Data | Differential privacy | Membership, inversion | Accuracy at low budgets | No neural curve |
| Data | Federated learning, secure aggregation | Raw data exposure | Communication | Demonstrated on EEG |
| Data | Homomorphic inference | Cloud exposure | High training cost | Demonstrated on EEG |
| Governance | Consent, provenance ledgers | Secondary use | Administrative | Proposed |
| Governance | Minimum reporting standards | Opacity | Administrative | Security not covered |

## 10. RESEARCH GAPS AND THE ROAD AHEAD

Figure 9 positions the open problems by research maturity against urgency for safe healthcare deployment, and proposes a staged agenda.

eISSN: 2583-1194

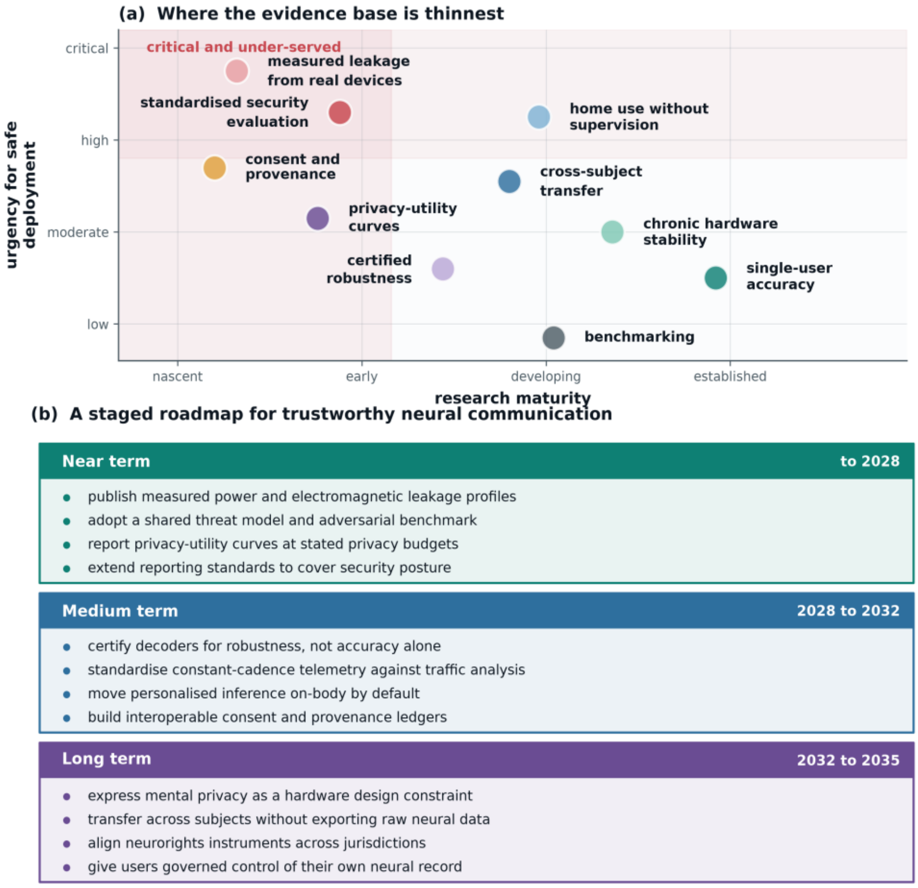


**Figure 9**. Research gap map and a staged roadmap to 2035.

Five gaps stand out, and four of them are security gaps.

No measured leakage from real devices. Power and electromagnetic analysis of embedded neural network inference is established [53, 10], and neural interfaces run on precisely that hardware, yet no published study reports a leakage profile from a fielded interface. Until such measurements exist, every claim about the exploitability of physical side channels here, including those organised in Figure 7, remains an inference from adjacent hardware.

No standard security evaluation. The field benchmarks accuracy rigorously through shared datasets and reproducibility frameworks [34], and does not benchmark security at all. There is no agreed threat model, no adversarial test suite reported alongside accuracy, and no requirement that a decoder disclose its robustness. The alignment-based defence benchmark is the closest existing artefact [67] and should become a reporting requirement.

No privacy-utility curve for neural data. The medical deep-learning community can state what a given privacy budget costs [62]. This domain cannot. Producing the equivalent curve for motor imagery, evoked-potential and speech decoding is tractable and high-value.

Generalisation and the participation gap. Cross-subject transfer remains the practical barrier to deployment, and foundation models are the current best answer [3]. Whether pretrained representations dissolve the participation gap identified two decades ago [31] or merely relocate it is unresolved.

Autonomous use outside supervision. Sustained independent home use is demonstrated [6], but the evidence base is a small number of participants. Chronic stability is quantified for one array type [24] and largely unknown for the newest high-density devices.

The staged roadmap in Figure 9(b) follows. The near-term actions are measurement and disclosure: publish leakage profiles, adopt a shared threat model and adversarial benchmark, report privacy-utility curves, and extend minimum reporting standards to security posture. The medium-term actions are architectural: certify decoders for robustness rather than accuracy alone, standardise constant-cadence telemetry so traffic shape carries no paradigm information, move personalised inference on-body by default, and build interoperable consent ledgers. The long-term actions are structural: express mental privacy as a hardware design constraint rather than a policy afterthought, deliver foundation models that transfer without exporting raw neural data, and align neurorights instruments internationally.

## 11. CONCLUSION

Across a hundred and fifty years, neural communication systems have passed through observation, control, restoration and, in the present decade, conversation, maturing into a genuine healthcare technology. Assistive throughput has risen from roughly two characters per minute in 1988 to a sustained fifty-six words per minute produced independently at home over two years, an improvement of more than two orders of magnitude, achieved through the compounding of high-density electrode arrays, self-supervised representation learning and language modelling rather than any single breakthrough.

The same convergence has produced a class of device that is wireless, always on, connected to remote inference and physically bound to a person, and whose security has not been characterised at the level the capability warrants. The classical side-channel literature demonstrates that correct implementations leak through power, electromagnetic, timing and traffic observables, and that embedded inference of exactly the kind these systems perform leaks both model structure and private inputs. The brain-specific literature demonstrates that stimuli can be weaponised to probe memory, that imperceptible perturbations can drive spellers to arbitrary output, and that neural recordings are strongly re-identifying. What is missing is not the method but its application.

The asymmetry is the finding. Decoding capability is ahead of the security evidence base by a wide margin, and the harms in question are not recoverable in the way a compromised credential is. Closing that gap requires measurement before regulation and disclosure before certification. The technical means exist in adjacent fields; this field needs to adopt them, and to treat the security posture of a neural communication system as a reportable property rather than an externality.

## REFERENCES

[1] J. R. Wolpaw, N. Birbaumer, D. J. McFarland, G. Pfurtscheller, and T. M. Vaughan, "Brain-computer interfaces for communication and control," Clinical Neurophysiology, vol. 113, no. 6, pp. 767-791, Jun. 2002, doi: 10.1016/S1388-2457(02)00057-3.

[2] Y. Roy, H. Banville, I. Albuquerque, A. Gramfort, T. H. Falk, and J. Faubert, "Deep learning-based electroencephalography analysis: a systematic review," Journal of Neural Engineering, vol. 16, no. 5, art. 051001, Aug. 2019, doi: 10.1088/1741-2552/ab260c.

[3] W.-B. Jiang, L.-M. Zhao, and B.-L. Lu, "Large brain model for learning generic representations with tremendous EEG data in BCI," in Proceedings of the Twelfth International Conference on Learning Representations, Vienna, Austria, 2024, arXiv:2405.18765.

[4] M. Hettick, E. Ho, A. J. Poole, M. Monge, D. Papageorgiou, K. Takahashi, M. LaMarca, D. Trietsch, K. Reed, M. Murphy, S. Rider, K. R. Gelman, Y. W. Byun, J. S. Miller, T. Hanson, V. Tolosa, S.-H. Lee, S. Bhatia, P. E. Konrad, M. Mager, C. H. Mermel, and B. I. Rapoport, "Minimally invasive implantation of scalable high-density cortical microelectrode arrays for multimodal neural decoding and stimulation," Nature Biomedical Engineering, vol. 10, no. 6, pp. 1206-1221, Oct. 2025, doi: 10.1038/s41551-025-01501-w.

[5] X. Zhang, L. Yao, S. Zhang, S. Kanhere, M. Sheng, and Y. Liu, "Internet of Things meets brain-computer interface: a unified deep learning framework for enabling human-thing cognitive interactivity," IEEE Internet of Things Journal, vol. 6, no. 2, pp. 2084-2092, Apr. 2019, doi: 10.1109/JIOT.2018.2877786.

[6] N. S. Card, T. Singer-Clark, H. Peracha, C. Iacobacci, X. Hou, M. Wairagkar, Z. Fogg, E. C. Offenberg, L. R. Hochberg, S. D. Stavisky, and D. M. Brandman, "Long-term independent use of an intracortical brain-computer interface for speech and cursor control," Nature Medicine, vol. 32, no. 7, pp. 2504-2510, 2026, doi: 10.1038/s41591-026-04414-6.

[7] I. Martinovic, D. Davies, M. Frank, D. Perito, T. Ros, and D. Song, "On the feasibility of side-channel attacks with brain-computer interfaces," in Proceedings of the 21st USENIX Security Symposium, Bellevue, WA, USA, Aug. 2012, pp. 143-158.

[8] T. Denning, Y. Matsuoka, and T. Kohno, "Neurosecurity: security and privacy for neural devices," Neurosurgical Focus, vol. 27, no. 1, art. E7, Jul. 2009, doi: 10.3171/2009.4.FOCUS0985.

[9] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in Advances in Cryptology - CRYPTO 1999, Lecture Notes in Computer Science, vol. 1666. Berlin, Germany: Springer, 1999, pp. 388-397, doi: 10.1007/3-540-48405-1_25.

[10] L. Batina, S. Bhasin, D. Jap, and S. Picek, "CSI NN: reverse engineering of neural network architectures through electromagnetic side channel," in Proceedings of the 28th USENIX Security Symposium, Santa Clara, CA, USA, Aug. 2019, pp. 515-532.

[11] E. Musk and Neuralink, "An integrated brain-machine interface platform with thousands of channels," Journal of Medical Internet Research, vol. 21, no. 10, art. e16194, Oct. 2019, doi: 10.2196/16194.

[12] H. Berger, "Uber das Elektrenkephalogramm des Menschen," Archiv fur Psychiatrie und Nervenkrankheiten, vol. 87, no. 1, pp. 527-570, 1929, doi: 10.1007/BF01797193.

[13] J. J. Vidal, "Toward direct brain-computer communication," Annual Review of Biophysics and Bioengineering, vol. 2, pp. 157-180, 1973, doi: 10.1146/annurev.bb.02.060173.001105.

[14] L. A. Farwell and E. Donchin, "Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials," Electroencephalography and Clinical Neurophysiology, vol. 70, no. 6, pp. 510-523, Dec. 1988, doi: 10.1016/0013-4694(88)90149-6.

[15] N. Birbaumer, N. Ghanayim, T. Hinterberger, I. Iversen, B. Kotchoubey, A. Kubler, J. Perelmouter, E. Taub, and H. Flor, "A spelling device for the paralysed," Nature, vol. 398, no. 6725, pp. 297-298, Mar. 1999, doi: 10.1038/18581.

[16] G. Pfurtscheller and F. H. Lopes da Silva, "Event-related EEG/MEG synchronization and desynchronization: basic principles," Clinical Neurophysiology, vol. 110, no. 11, pp. 1842-1857, Nov. 1999, doi: 10.1016/S1388-2457(99)00141-8.

[17] E. C. Leuthardt, G. Schalk, J. R. Wolpaw, J. G. Ojemann, and D. W. Moran, "A brain-computer interface using electrocorticographic signals in humans," Journal of Neural Engineering, vol. 1, no. 2, pp. 63-71, Jun. 2004, doi: 10.1088/1741-2560/1/2/001.

[18] L. R. Hochberg, M. D. Serruya, G. M. Friehs, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue, "Neuronal ensemble control of prosthetic devices by a human with tetraplegia," Nature, vol. 442, no. 7099, pp. 164-171, Jul. 2006, doi: 10.1038/nature04970.

[19] L. R. Hochberg, D. Bacher, B. Jarosiewicz, N. Y. Masse, J. D. Simeral, J. Vogel, S. Haddadin, J. Liu, S. S. Cash, P. van der Smagt, and J. P. Donoghue, "Reach and grasp by people with tetraplegia using a neurally controlled robotic arm," Nature, vol. 485, no. 7398, pp. 372-375, May 2012, doi: 10.1038/nature11076.

[20] H. Ramoser, J. Muller-Gerking, and G. Pfurtscheller, "Optimal spatial filtering of single trial EEG during imagined hand movement," IEEE Transactions on Rehabilitation Engineering, vol. 8, no. 4, pp. 441-446, Dec. 2000, doi: 10.1109/86.895946.

[21] K. K. Ang, Z. Y. Chin, C. Wang, C. Guan, and H. Zhang, "Filter bank common spatial pattern algorithm on BCI Competition IV Datasets 2a and 2b," Frontiers in Neuroscience, vol. 6, art. 39, Mar. 2012, doi: 10.3389/fnins.2012.00039.

[22] A. Barachant, S. Bonnet, M. Congedo, and C. Jutten, "Multiclass brain-computer interface classification by Riemannian geometry," IEEE Transactions on Biomedical Engineering, vol. 59, no. 4, pp. 920-928, Apr. 2012, doi: 10.1109/TBME.2011.2172210.

[23] X. Chen, Y. Wang, M. Nakanishi, X. Gao, T.-P. Jung, and S. Gao, "High-speed spelling with a noninvasive brain-computer interface," Proceedings of the National Academy of Sciences of the United States of America, vol. 112, no. 44, pp. E6058-E6067, Nov. 2015, doi: 10.1073/pnas.1508080112.

[24] C. Sponheim, V. Papadourakis, J. L. Collinger, J. Downey, J. Weiss, L. Pentousi, K. Elliott, and N. G. Hatsopoulos, "Longevity and reliability of chronic unit recordings using the Utah, intracortical multi-electrode arrays," Journal of Neural Engineering, vol. 18, no. 6, art. 066044, Dec. 2021, doi: 10.1088/1741-2552/ac3eaf.

[25] N. A. Steinmetz, C. Aydin, A. Lebedeva, M. Okun, and M. Pachitariu, "Neuropixels 2.0: a miniaturized high-density probe for stable, long-term brain recordings," Science, vol. 372, no. 6539, art. eabf4588, Apr. 2021, doi: 10.1126/science.abf4588.

[26] P. Mitchell, S. C. M. Lee, P. E. Yoo, A. Morokoff, R. P. Sharma, D. L. Williams, C. MacIsaac, M. E. Howard, L. Irving, I. Vrljic, C. Williams, S. Bush, A. H. Balabanski, K. J. Drummond, P. Desmond, D. Weber, T. Denison, S. Mathers, T. J. O'Brien, J. Mocco, D. B. Grayden, D. S. Liebeskind, N. L. Opie, T. J. Oxley, and B. C. V. Campbell, "Assessment of safety of a fully implanted endovascular brain-computer interface for severe paralysis in 4 patients: the Stentrode with Thought-Controlled Digital Switch (SWITCH) study," JAMA Neurology, vol. 80, no. 3, pp. 270-278, Mar. 2023, doi: 10.1001/jamaneurol.2022.4847.

[27] N. Kaongoen, J. Choi, J. W. Choi, H. Kwon, C. Hwang, G. Hwang, B. H. Kim, and S. Jo, "The future of wearable EEG: a review of ear-EEG technology and its applications," Journal of Neural Engineering, vol. 20, no. 5, art. 051002, Oct. 2023, doi: 10.1088/1741-2552/acfcda.

[28] E. Boto, N. Holmes, J. Leggett, G. Roberts, V. Shah, S. S. Meyer, L. Duque Munoz, K. J. Mullinger, T. M. Tierney, S. Bestmann, G. R. Barnes, R. Bowtell, and M. J. Brookes, "Moving magnetoencephalography towards real-world applications with a wearable system," Nature, vol. 555, no. 7698, pp. 657-661, Mar. 2018, doi: 10.1038/nature26147.

[29] J. Sabio, N. S. Williams, G. M. McArthur, and N. A. Badcock, "A scoping review on the use of consumer-grade EEG devices for research," PLOS ONE, vol. 19, no. 3, art. e0291186, Mar. 2024, doi: 10.1371/journal.pone.0291186.

[30] F. Lotte, L. Bougrain, A. Cichocki, M. Clerc, M. Congedo, A. Rakotomamonjy, and F. Yger, "A review of classification algorithms for EEG-based brain-computer interfaces: a 10 year update," Journal of Neural Engineering, vol. 15, no. 3, art. 031005, Apr. 2018, doi: 10.1088/1741-2552/aab2f2.

[31] C. Vidaurre and B. Blankertz, "Towards a cure for BCI illiteracy," Brain Topography, vol. 23, no. 2, pp. 194-198, Jun. 2010, doi: 10.1007/s10548-009-0121-6.

[32] R. T. Schirrmeister, J. T. Springenberg, L. D. J. Fiederer, M. Glasstetter, K. Eggensperger, M. Tangermann, F. Hutter, W. Burgard, and T. Ball, "Deep learning with convolutional neural networks for EEG decoding and visualization," Human Brain Mapping, vol. 38, no. 11, pp. 5391-5420, Nov. 2017, doi: 10.1002/hbm.23730.

[33] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces," Journal of Neural Engineering, vol. 15, no. 5, art. 056013, Jul. 2018, doi: 10.1088/1741-2552/aace8c.

[34] V. Jayaram and A. Barachant, "MOABB: trustworthy algorithm benchmarking for BCIs," Journal of Neural Engineering, vol. 15, no. 6, art. 066011, Sep. 2018, doi: 10.1088/1741-2552/aadea0.

[35] D. Kostas, S. Aroca-Ouellette, and F. Rudzicz, "BENDR: using transformers and a contrastive self-supervised learning task to learn from massive amounts of EEG data," Frontiers in Human Neuroscience, vol. 15, art. 653659, 2021, doi: 10.3389/fnhum.2021.653659.

[36] J. Tang, A. LeBel, S. Jain, and A. G. Huth, "Semantic reconstruction of continuous language from non-invasive brain recordings," Nature Neuroscience, vol. 26, no. 5, pp. 858-866, May 2023, doi: 10.1038/s41593-023-01304-9.

[37] A. Defossez, C. Caucheteux, J. Rapin, O. Kabeli, and J.-R. King, "Decoding speech perception from non-invasive brain recordings," Nature Machine Intelligence, vol. 5, pp. 1097-1107, 2023, doi: 10.1038/s42256-023-00714-5.

[38] G. K. Anumanchipalli, J. Chartier, and E. F. Chang, "Speech synthesis from neural decoding of spoken sentences," Nature, vol. 568, no. 7753, pp. 493-498, Apr. 2019, doi: 10.1038/s41586-019-1119-1.

[39] D. A. Moses, S. L. Metzger, J. R. Liu, G. K. Anumanchipalli, J. G. Makin, P. F. Sun, J. Chartier, M. E. Dougherty, P. M. Liu, G. M. Abrams, A. Tu-Chan, K. Ganguly, and E. F. Chang, "Neuroprosthesis for decoding speech in a paralyzed person with anarthria," The New England Journal of Medicine, vol. 385, no. 3, pp. 217-227, Jul. 2021, doi: 10.1056/NEJMoa2027540.

[40] F. R. Willett, D. T. Avansino, L. R. Hochberg, J. M. Henderson, and K. V. Shenoy, "High-performance brain-to-text communication via handwriting," Nature, vol. 593, no. 7858, pp. 249-254, May 2021, doi: 10.1038/s41586-021-03506-2.

[41] F. R. Willett, E. M. Kunz, C. Fan, D. T. Avansino, G. H. Wilson, E. Y. Choi, F. Kamdar, M. F. Glasser, L. R. Hochberg, S. Druckmann, K. V. Shenoy, and J. M. Henderson, "A high-performance speech neuroprosthesis," Nature, vol. 620, no. 7976, pp. 1031-1036, Aug. 2023, doi: 10.1038/s41586-023-06377-x.

[42] S. L. Metzger, K. T. Littlejohn, A. B. Silva, D. A. Moses, M. P. Seaton, R. Wang, M. E. Dougherty, J. R. Liu, P. Wu, M. A. Berger, I. Zhuravleva, A. Tu-Chan, K. Ganguly, G. K. Anumanchipalli, and E. F. Chang, "A high-performance neuroprosthesis for speech decoding and avatar control," Nature, vol. 620, no. 7976, pp. 1037-1046, Aug. 2023, doi: 10.1038/s41586-023-06443-4.

[43] N. S. Card, M. Wairagkar, C. Iacobacci, X. Hou, T. Singer-Clark, F. R. Willett, E. M. Kunz, C. Fan, M. Vahdati Nia, D. R. Deo, A. Srinivasan, E. Y. Choi, M. F. Glasser, L. R. Hochberg, J. M. Henderson, K. Shahlaie, S. D. Stavisky, and D. M. Brandman, "An accurate and rapidly calibrating speech neuroprosthesis," The New England Journal of Medicine, vol. 391, no. 7, pp. 609-618, Aug. 2024, doi: 10.1056/NEJMoa2314132.

[44] M. Wairagkar, N. S. Card, T. Singer-Clark, X. Hou, C. Iacobacci, L. M. Miller, L. R. Hochberg, D. M. Brandman, and S. D. Stavisky, "An instantaneous voice-synthesis neuroprosthesis," Nature, vol. 644, no. 8075, pp. 145-152, Aug. 2025, doi: 10.1038/s41586-025-09127-3.

[45] J. Levy, M. Zhang, S. Pinet, J. Rapin, H. Banville, S. d'Ascoli, and J.-R. King, "Brain-to-text decoding: a non-invasive approach via typing," arXiv:2502.17480, Feb. 2025.

[46] N. C. Camgoz, O. Koller, S. Hadfield, and R. Bowden, "Sign language transformers: joint end-to-end sign language recognition and translation," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, Jun. 2020, pp. 10020-10030, doi: 10.1109/CVPR42600.2020.01004.

[47] Defense Advanced Research Projects Agency, "Next-Generation Nonsurgical Neurotechnology (N3)," DARPA, Arlington, VA, USA, programme record. [Online]. Available: https://www.darpa.mil/research/programs/next-generation-nonsurgical-neurotechnology (accessed Jul. 25, 2026).

[48] A. Binnendijk, T. Marler, and E. M. Bartels, Brain-Computer Interfaces: U.S. Military Applications and Implications, An Initial Assessment, Report RR-2996-RC. Santa Monica, CA, USA: RAND Corporation, 2020, doi: 10.7249/RR2996.

[49] K. LaFleur, K. Cassady, A. Doud, K. Shades, E. Rogin, and B. He, "Quadcopter control in three-dimensional space using a noninvasive motor imagery-based brain-computer interface," Journal of Neural Engineering, vol. 10, no. 4, art. 046003, Aug. 2013, doi: 10.1088/1741-2560/10/4/046003.

[50] K. A. Feltman, A. M. Hayes, K. A. Bernhardt, E. Nwala, and A. M. Kelley, "Viability of tDCS in military environments for performance enhancement: a systematic review," Military Medicine, vol. 185, no. 1-2, pp. e53-e60, Jan. 2020, doi: 10.1093/milmed/usz189.

[51] X. Wang, M. Hersche, M. Magno, and L. Benini, "MI-BMInet: an efficient convolutional neural network for motor imagery brain-machine interfaces with EEG channel selection," IEEE Sensors Journal, vol. 24, no. 6, pp. 8835-8847, Mar. 2024, doi: 10.1109/JSEN.2024.3353146.

[52] P. Sirinam, M. Imani, M. Juarez, and M. Wright, "Deep fingerprinting: undermining website fingerprinting defenses with deep learning," in Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, Toronto, Canada, Oct. 2018, pp. 1928-1943, doi: 10.1145/3243734.3243768.

[53] L. Wei, B. Luo, Y. Li, Y. Liu, and Q. Xu, "I know what you see: power side-channel attack on convolutional neural network accelerators," in Proceedings of the 34th Annual Computer Security Applications Conference, San Juan, PR, USA, Dec. 2018, pp. 393-406, doi: 10.1145/3274694.3274696.

[54] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in Proceedings of the 2017 IEEE Symposium on Security and Privacy, San Jose, CA, USA, May 2017, pp. 3-18, doi: 10.1109/SP.2017.41.

[55] L. Pycroft, S. G. Boccard, S. L. F. Owen, J. F. Stein, J. J. Fitzgerald, A. L. Green, and T. Z. Aziz, "Brainjacking: implant security issues in invasive neuromodulation," World Neurosurgery, vol. 92, pp. 454-462, Aug. 2016, doi: 10.1016/j.wneu.2016.05.010.

[56] X. Zhang and D. Wu, "On the vulnerability of CNN classifiers in EEG-based BCIs," IEEE Transactions on Neural Systems and Rehabilitation Engineering, vol. 27, no. 5, pp. 814-825, May 2019, doi: 10.1109/TNSRE.2019.2908955.

[57] X. Zhang, D. Wu, L. Ding, H. Luo, C.-T. Lin, T.-P. Jung, and R. Chavarriaga, "Tiny noise, big mistakes: adversarial perturbations induce errors in brain-computer interface spellers," National Science Review, vol. 8, no. 4, art. nwaa233, Apr. 2021, doi: 10.1093/nsr/nwaa233.

[58] L. Meng, X. Jiang, J. Huang, W. Li, H. Luo, and D. Wu, "EEG-based brain-computer interfaces are vulnerable to backdoor attacks," IEEE Transactions on Neural Systems and Rehabilitation Engineering, vol. 31, pp. 2224-2234, 2023, doi: 10.1109/TNSRE.2023.3273214.

[59] M. V. Ruiz-Blondet, Z. Jin, and S. Laszlo, "CEREBRE: a novel method for very high accuracy event-related potential biometric identification," IEEE Transactions on Information Forensics and Security, vol. 11, no. 7, pp. 1618-1629, Jul. 2016, doi: 10.1109/TIFS.2016.2543524.

[60] Y. Ishai, A. Sahai, and D. Wagner, "Private circuits: securing hardware against probing attacks," in Advances in Cryptology - CRYPTO 2003, Lecture Notes in Computer Science, vol. 2729. Berlin, Germany: Springer, 2003, pp. 463-481, doi: 10.1007/978-3-540-45146-4_27.

[61] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, Oct. 2016, pp. 308-318, doi: 10.1145/2976749.2978318.

[62] M. Mohammadi, M. Vejdanihemmat, M. Lotfinia, M. Rusu, D. Truhn, A. Maier, and S. Tayebi Arasteh, "Differential privacy for medical deep learning: methods, tradeoffs, and deployment implications," npj Digital Medicine, vol. 9, art. 93, 2026, doi: 10.1038/s41746-025-02280-z.

[63] T. Jia, L. Meng, S. Li, J. Liu, and D. Wu, "Federated motor imagery classification for privacy-preserving brain-computer interfaces," IEEE Transactions on Neural Systems and Rehabilitation Engineering, vol. 32, pp. 3442-3451, 2024, doi: 10.1109/TNSRE.2024.3457504.

[64] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," in Advances in Neural Information Processing Systems 32, Vancouver, Canada, Dec. 2019, doi: 10.48550/arXiv.1906.08935.

[65] A. B. Popescu, I. A. Taca, C. I. Nita, A. Vizitiu, R. Demeter, C. Suciu, and L. M. Itu, "Privacy preserving classification of EEG data using machine learning and homomorphic encryption," Applied Sciences, vol. 11, no. 16, art. 7360, 2021, doi: 10.3390/app11167360.

[66] X. Chen, Z. Wang, and D. Wu, "Alignment-based adversarial training (ABAT) for improving the robustness and accuracy of EEG-based BCIs," IEEE Transactions on Neural Systems and Rehabilitation Engineering, vol. 32, pp. 1703-1714, 2024, doi: 10.1109/TNSRE.2024.3391936.

[67] X. Chen, T. Jia, and D. Wu, "Data alignment based adversarial defense benchmark for EEG-based BCIs," Neural Networks, vol. 188, art. 107516, Aug. 2025, doi: 10.1016/j.neunet.2025.107516.

[68] S. Lopez Bernal, A. Huertas Celdran, G. Martinez Perez, M. T. Barros, and S. Balasubramaniam, "Security in brain-computer interfaces: state-of-the-art, opportunities, and future challenges," ACM Computing Surveys, vol. 54, no. 1, art. 11, pp. 1-35, Jan. 2021, doi: 10.1145/3427376.

[69] M. Ienca and R. Andorno, "Towards new human rights in the age of neuroscience and neurotechnology," Life Sciences, Society and Policy, vol. 13, no. 1, art. 5, Apr. 2017, doi: 10.1186/s40504-017-0050-1.

[70] C. Eiber, J. Delbeke, J. Cardoso, M. de Neeling, S. John, C. W. Lee, J. Skefos, A. Sun, D. Prodanov, and Z. McKinney, "Preliminary minimum reporting requirements for in-vivo neural interface research: I. implantable neural interfaces," IEEE Open Journal of Engineering in Medicine and Biology, vol. 2, pp. 74-83, Feb. 2021, doi: 10.1109/OJEMB.2021.3060919.