Leadership, Learning, and Business Beyond Boundaries: Managing Change in a Complex World
21–23 May 2026
Ljubljana, Slovenia
make learn TIIM PIConf
Management, Knowledge and Learning International Conference
Technology, Innovation and Industrial Management
Sustainable Development

Ethics and Trust in Brain-Computer Interfaces: A Theoretical Review

Gregor Jagodič
MFDPŠ Celje, Slovenia
gregor.jagodic@mfdps.si

Rahela Aničić Čandrlić
PhD Candidate at DOBA Business School Maribor, Slovenia
rahela.anicic.phd@gmail.com

Abstract

Brain-computer interfaces (BCIs) are increasingly developed for therapeutic, enhancement-oriented, and consumer applications, making ethical reflection central to their responsible development and use. This theoretical review examines contemporary BCI ethics through a structured narrative review of 15 peer-reviewed publications from 2018 to 2025, supplemented by methodological and contextual sources. Its distinctive contribution is a context-sensitive synthesis showing that ethical concerns do not carry the same weight across all applications. Therapeutic BCIs foreground vulnerability, informed consent, agency, and continuity of care; enhancement-oriented uses intensify concerns about fairness, authenticity, and social pressure; while commercial applications raise risks of hype, weak evidence, data commodification, and misplaced trust. The article argues that trust is a practical indicator of whether privacy, user control, accountability, and justice are adequately addressed. It concludes that ethically acceptable BCI development requires human-centred governance, transparent accountability, meaningful user control, and context-sensitive regulation to support legitimacy and sustainable use.

Keywords: brain-computer interfaces (BCI); neuroethics; trust; governance; neurodata

INTRODUCTION

Brain-computer interfaces (BCIs) have become one of the most important developments in contemporary neurotechnology. By creating a direct pathway between neural activity and external computational systems, BCIs enable the translation of brain signals into commands, communication, or feedback. In clinical contexts, they are primarily associated with therapeutic aims, such as restoring communication, supporting motor recovery, and assisting individuals with severe neurological impairments. However, their significance extends beyond medicine. Recent scholarship increasingly recognises that BCIs are socio-technical systems that shape relationships among human agency,

machine interpretation, institutional oversight, and commercial interests (Cabrera & Weber, 2023; Poo, 2024).

BCIs are ethically distinctive because they operate at the intersection of neural processes and digital infrastructures. Unlike many other emerging technologies, they engage with signals closely connected to intention, cognition, communication, and aspects of personal identity. Ethical analysis, therefore, cannot be limited to technical performance or medical safety alone. It must also address broader issues such as privacy, autonomy, dignity, fairness, and governance, particularly as BCIs move from experimental settings into clinical and commercial practice (Yuste, 2023; Cabrera & Weber, 2023).

The diversification of their applications further reinforces the ethical relevance of BCIs. While therapeutic uses remain central, enhancement-oriented and consumer-facing applications are becoming more visible. This shift is important because a BCI designed to restore lost function is not evaluated according to the same moral criteria as a device intended for cognitive enhancement, workplace monitoring, entertainment, or wellness. The acceptable level of risk, the justification for use, and the scope of regulation all vary across contexts. Consequently, ethical assessment must remain sensitive to the purposes, settings, and social conditions in which BCIs are developed and used (Wexler & Thibault, 2019; Gordon & Seth, 2024).

The distinctive contribution of this article is not the identification of entirely new ethical categories, but the development of a context-sensitive synthesis that links ethical domains to different forms of BCI use. Existing scholarship often discusses privacy, autonomy, identity, justice, trust, and governance as separate concerns. This article argues that these domains should instead be understood as mutually dependent components of BCI legitimacy. The analysis, therefore, contributes a conceptual framework in which the ethical acceptability of BCIs depends on the interaction between technological design, context of use, user vulnerability, data governance, and institutional accountability.

Statement of the Research Problem

The central research problem addressed in this article stems from the growing gap between the rapid development of brain-computer interfaces and the slower development of ethical frameworks to guide their use. Recent advances in machine learning, neural signal decoding, implantable systems, and speech-generating interfaces have strengthened expectations regarding the transformative potential of BCIs. At the same time, ethical reflection has often remained fragmented, focusing mainly on safety, feasibility, and immediate risk-benefit considerations. Although these issues remain important, they no longer capture the full ethical complexity of contemporary BCI development (Specker Sullivan & Illes, 2018; Cabrera & Weber, 2023).

An additional challenge lies in the tendency to treat BCIs as a single technological category with uniform ethical implications. BCI systems differ in their invasiveness, purpose, technical architecture, and context of use. Some are designed to restore communication or motor function in vulnerable patients, whereas others are being explored for enhancement, consumer applications, or commercial digital services. As a result, ethical concerns vary across contexts and include not only patient vulnerability and long-term support, but also data commodification, exaggerated market claims, and reduced user control in AI-assisted systems (Wexler & Thibault, 2019; van Stuijvenberg et al., 2024; Boonstra, 2025).

The increasing ethical sensitivity of neural data also shapes the research problem. Recent scholarship emphasises that neurodata may reveal information about attention, intention, affective states, and communication, making privacy, informed consent, and data governance central rather than secondary concerns. At the same time, BCIs raise broader questions about autonomy, agency, and identity, particularly in systems where machine interpretation influences communication and action. The issue is therefore not whether BCIs are simply beneficial or harmful technologies, but which ethical domains are most central in current scholarship, how they change across contexts of use, and which governance principles are necessary to support ethically responsible BCI development (Yuste, 2023; Poo, 2024; Gordon & Seth, 2024).

Aim and Objectives of the Study

This article aims to provide a theoretically grounded analysis of the ethical challenges associated with brain-computer interfaces in contemporary academic scholarship. More specifically, it seeks to clarify the conceptual structure of BCI ethics by identifying the ethical domains that recur most consistently in recent literature and examining their relationships. Rather than treating privacy, autonomy, identity, justice, governance, and trust as isolated themes, the article approaches them as interconnected dimensions of a broader ethical field. This perspective is important because the ethical legitimacy of BCI systems depends not on any single factor in isolation, but on the interaction of multiple normative considerations within concrete technological and institutional settings (Cabrera & Weber, 2023; Yuste, 2023).

The article also seeks to distinguish between different contexts of BCI application. Recent scholarship suggests that ethical analysis becomes more precise when therapeutic, enhancement-oriented, and commercial uses are not treated as a single evaluative category. Accordingly, the discussion examines how issues such as vulnerability, voluntariness, fairness, accountability, and user control shift depending on whether BCIs are used to restore lost function, augment existing capacities, or support consumer-oriented and commercial applications (Wexler & Thibault, 2019; Gordon & Seth, 2024; Boonstra, 2025).

In addition, the article identifies governance principles that can support the ethically acceptable development and implementation of BCIs. In recent scholarship, ethical oversight is increasingly understood as something that should be integrated into design, clinical translation, data governance, and regulation from the outset. For this reason, the analysis synthesises the principles that appear most robust across contemporary literature, including meaningful consent, neurodata protection, transparent accountability, user control, and fairness in access and benefit distribution (Poo, 2024; Yuste, 2023; van Stuijvenberg et al., 2024).

Overall, the article is guided by the assumption that ethics should not be treated as an external constraint on innovation. In the case of BCIs, ethics functions as a condition of legitimacy, since the acceptability of these technologies depends not only on technical effectiveness, but also on whether they respect users, protect vulnerable populations, preserve agency, and distribute risks and benefits in a justifiable manner (Cabrera & Weber, 2023; Gordon & Seth, 2024).

Research Questions

In line with the research problem and the stated aims of the study, the article is guided by three research questions:

- RQ1: Which ethical domains are most central in contemporary academic discussions of BCI?
- RQ2: How do ethical concerns related to BCI change across therapeutic, enhancement-oriented, and commercial contexts of use?
- RQ3: Which governance principles emerge from recent scholarship as necessary for the ethically acceptable development and implementation of BCI?

These research questions follow a funnel-shaped analytical logic from general to more specific issues. The first question identifies the main ethical domains that structure current BCI scholarship. The second examines how these concerns vary across different contexts of use. The third moves toward normative synthesis by identifying the governance principles that can help align BCI innovation with broader standards of autonomy, justice, privacy, accountability, and trust. Together, these questions provide a coherent structure for the remainder of the article and support a focused and theoretically integrated discussion (van Stuijvenberg et al., 2024; Boonstra, 2025).

THEORETICAL FRAMEWORK

Definition and Characteristics of BCI Technology

Brain-computer interface technology is generally defined as a system that acquires neural signals, translates them into machine-readable information, and uses those signals to enable communication, control, or intervention without relying exclusively on conventional muscular pathways. In practical terms, a BCI translates brain activity into an external device, such as a computer cursor, a speech synthesiser, a robotic system, or a stimulation platform. However, this definition captures only the technology's technical dimension. From a theoretical perspective, BCIs must also be understood as socio-technical systems located at the intersection of neuroscience, medicine, artificial intelligence, data processing, and human experience (Cabrera & Weber, 2023; Poo, 2024).

A key characteristic of BCIs is that neural activity becomes meaningful only through computational mediation. Processes such as signal acquisition, preprocessing, decoding, and output generation shape how brain activity is translated into communication or action. As a result, BCIs do not simply “read” the brain directly, but construct operational interpretations of neural signals within specific technical and contextual frameworks. This means that BCI outputs are influenced not only by the user’s neural states, but also by algorithms, device design, training data, and system architecture (van Stuijvenberg et al., 2024; Soldado-Magraner et al., 2024).

BCIs also vary in their technical form and degree of invasiveness. Some systems are invasive and require implanted electrodes, while others are partially invasive or non-invasive. This distinction has important implications for risk, reliability, long-term dependence, and the balance between burden and benefit. Invasive systems often provide greater precision, particularly in advanced therapeutic applications, whereas non-invasive systems are generally more accessible but less stable and accurate (Poo, 2024; Soldado-Magraner et al., 2024).

Another important feature of contemporary BCIs is their growing integration with adaptive and, in some cases, bidirectional systems. Current developments increasingly include closed-loop arrangements and

AI-assisted processes of prediction and output generation. While this expands the practical potential of BCIs, it also introduces new forms of opacity, since the relation between user intention and system response may become more heavily mediated by algorithms. For this reason, recent scholarship treats BCIs not merely as assistive tools, but as neuroethical objects that reshape agency, control, and institutional responsibility (Haag et al., 2025; van Stuijvenberg et al., 2024).

Fields of Application of BCI Technology

The application landscape of BCI technology has expanded considerably in recent years, intensifying both theoretical and ethical debate. Historically, BCIs were primarily associated with assistive and therapeutic purposes, especially for individuals with profound motor impairment. In this context, they were understood mainly as restorative technologies aimed at recovering lost capacities rather than enhancing existing ones. Contemporary developments, however, show that BCI applications now extend well beyond this original clinical focus (Cabrera & Weber, 2023; Gordon & Seth, 2024).

One of the most established areas of application is clinical communication and motor assistance. BCIs have been developed to support patients with severe paralysis, amyotrophic lateral sclerosis, spinal cord injury, stroke, and related conditions that impair speech or movement. Their value lies in restoring communication, enabling environmental control, and supporting participation in daily and social life. At the same time, the literature emphasises that therapeutic value depends not only on proof of concept, but also on long-term usability, technical stability, and continued support beyond experimental settings (Kögel et al., 2020; Baker et al., 2023).

Another important area concerns neurorehabilitation and restorative intervention. In these cases, BCIs are used not only for communication or control but also to support recovery processes after neurological damage, including motor rehabilitation, adaptive stimulation, and closed-loop neuromodulation. Because such systems are more deeply integrated into treatment, they also raise questions about bodily integrity, dependence, and long-term clinical responsibility (Baker et al., 2023; Haag et al., 2025).

An increasingly visible application area is speech-related and communicative restoration. Recent work on speech-BCIs shows the growing potential to decode intended speech, generate text, produce synthetic voice output, and support avatar-mediated communication for individuals who have lost natural speech capacities (Metzger et al., 2023; Wairagkar et al., 2025). These applications are ethically significant because they relate directly to self-expression, communication rights, and social recognition, while also raising concerns about semantic accuracy, authorship, emotional expression, and user control (van Stuijvenberg et al., 2024).

Beyond medicine, BCIs are increasingly discussed in relation to cognitive enhancement, performance optimisation, entertainment, and consumer-facing digital services. In these contexts, BCIs are no longer framed primarily as restorative tools, but as technologies intended to augment attention, efficiency, or user experience. This shift raises concerns about authenticity, social pressure, fairness, hype, weak evidentiary standards, and data commodification (Wexler & Thibault, 2019; Gordon & Seth, 2024; Boonstra, 2025).

Taken together, these application areas show that BCI technology cannot be understood as a single moral or practical category. Its significance depends on whether it is used to restore lost capacities, support rehabilitation, enable communication, or enhance performance in commercial settings. This diversity is

ethically important because it shapes how benefits, risks, and responsibilities should be interpreted (Cabrera & Weber, 2023; Gordon & Seth, 2024).

Illustrative Cases of BCI Ethics in Practice

Although this article is theoretical in design, recent BCI developments illustrate why ethical analysis cannot remain purely abstract. Speech-related BCIs provide a particularly important example. Recent neuroprosthetic systems have demonstrated the ability to translate neural activity into text, synthesised speech, voice output, and avatar-based communicative expression for individuals with severe paralysis or loss of speech (Metzger et al., 2023; Wairagkar et al., 2025). These developments show the restorative potential of BCIs, but they also make ethical questions more concrete. If a system generates speech on behalf of a user, ethical evaluation must address not only decoding accuracy, but also authorship, correction, timing, emotional expression, and the user's ability to reject or modify outputs that do not reflect their intention.

Implanted and long-term BCI systems raise a related but distinct set of ethical concerns. These technologies may offer substantial benefits for people with severe motor impairment, but they also create dependence on technical maintenance, software updates, clinical follow-up, and institutional support. In such cases, informed consent cannot be treated as a one-time research form. It must include a realistic discussion of uncertainty, long-term support, possible device failure, data use, and discontinuation. The ethical issue is therefore not only whether implantation is initially justified, but also whether responsibility is sustained throughout the device's lifecycle.

Consumer-facing and commercial neurotechnology make the practical relevance of BCI ethics visible in another way. When BCI-related devices are marketed for wellness, productivity, entertainment, or cognitive enhancement, users may encounter strong claims before the evidence base is mature. This creates a different ethical problem from therapeutic use: users are positioned not only as beneficiaries, but also as consumers and potential sources of valuable neural data. In this setting, trust can become vulnerable to branding, hype, and information asymmetries rather than being grounded in transparent evidence and accountable governance.

Ethical Aspects of the Use of BCI Technology

The ethical analysis of BCI technology has expanded from a narrow focus on safety and feasibility to a broader framework that includes privacy, autonomy, identity, justice, and governance. This shift reflects the recognition that BCIs are not only biomedical devices, but also systems that shape communication, agency, data practices, and social power. Ethical evaluation, therefore, extends beyond technical effectiveness and safety to questions of respect for users, protection of sensitive data, fair distribution of benefits and burdens, and legitimate oversight and accountability (Specker Sullivan & Illes, 2018; Cabrera & Weber, 2023).

One of the most frequently discussed issues is mental privacy and the governance of neurodata. Neural data are considered ethically distinctive because they are closely linked to attention, intention, affective processing, and communication. Although current BCI systems do not provide unrestricted access to inner life, the literature remains cautious because the interpretive power of such data may increase through artificial intelligence, multimodal integration, and large-scale aggregation. Ethical concern,

therefore, extends beyond unauthorised access to include storage, secondary use, cross-platform sharing, commercialisation, and long-term control of brain-derived information (Yuste, 2023; Poo, 2024).

Another major concern relates to autonomy, informed consent, and meaningful user control. In BCI contexts, autonomy cannot be reduced to a one-time act of consent; it must be understood as an ongoing condition shaped by comprehension, voluntariness, ongoing oversight, and the user's ability to influence how the system operates. This is especially important when BCIs become integrated into everyday communication or bodily control. Recent scholarship, therefore, emphasises dynamic consent, context-sensitive communication, and design strategies that preserve meaningful control over outputs and decision pathways (Kögel et al., 2020; van Stuijvenberg et al., 2024).

BCIs also raise important questions about identity, agency, and self-experience. In long-term or high-dependence applications, the interface may become part of how a person acts, communicates, or experiences control. This may be empowering when the technology restores lost capacities, but it may also be destabilising if outputs become unpredictable or if device removal disrupts an established sense of agency. Ethical reflection on BCIs must therefore address not only regulation and safety, but also the user's lived experience of authorship, embodiment, and selfhood (Gilbert et al., 2023; Kögel et al., 2020).

Justice, access, and commercialisation form another central ethical cluster. As BCI development increasingly attracts commercial investment, the literature warns that innovation may privilege market opportunity over social need. This raises concerns about unequal access, exaggerated public expectations, and the extraction of value from neural data without adequate protection or benefit-sharing. In enhancement contexts, additional concerns arise about unequal advantage, coercive expectations of optimisation, and the possible normalisation of neurotechnology as a condition of productivity or competitiveness (Pfotenhauer et al., 2021; Boonstra, 2025; Gordon & Seth, 2024).

Finally, ethical analysis must also address responsibility and governance. Because BCIs span users, clinicians, engineers, companies, and regulators, accountability can easily become diffuse. Recent scholarship, therefore, suggests that ethically acceptable innovation requires governance frameworks capable of clarifying roles, oversight obligations, and standards for long-term care. Additional work on the medical applications of BCIs also emphasises the need for differentiated ethical standards, particularly regarding implantable versus non-implantable systems, standardised procedures, and context-sensitive evaluation in clinical use (Soldado-Magraner et al., 2024; Haag et al., 2025; Zhang et al., 2024).

Trust in BCI Technology and Its Use

Trust has become an increasingly important concept in discussions of BCI technology because technical capability alone is not sufficient to secure acceptance, legitimacy, or sustained use. A BCI may demonstrate therapeutic promise or strong decoding performance, yet still fail to gain durable support if users, clinicians, and the broader public do not regard it as reliable, understandable, and normatively responsible. In this sense, trust should not be understood as a vague psychological preference but as a relational condition that links technical performance to perceived competence, safety, integrity, and accountability (Starke & Ienca, 2025).

In the context of BCIs, trust operates on several interconnected levels. At the user level, it concerns whether a person can rely on the system for communication, action, or therapeutic support without fearing loss of control, misinterpretation, or misuse of data. At the clinical level, it concerns whether professionals can rely on the technology's safety, reproducibility, and evidentiary basis. At the institutional and societal level, it concerns whether developers, regulators, and commercial actors are perceived as acting transparently and responsibly. These levels are closely related, since trust at one level can be weakened by failures at another (Pfotenhauer et al., 2021; Starke & Ienca, 2025).

Trust is particularly important in AI-driven and speech-related BCIs. In these systems, users may depend on algorithmically mediated interpretation for highly personal forms of communication. Trust, therefore, depends not only on whether the technology functions but also on whether users retain sufficient influence over timing, correction, semantic content, and expressive authorship. Recent scholarship also shows that speech-BCIs raise concerns about speech ownership, since AI-assisted communication may blur the boundaries between user intention and system-generated output. For this reason, trust is increasingly treated as a key condition of meaningful and sustainable BCI use rather than as a secondary issue that follows once technical performance has been established (van Stuijvenberg et al., 2024; Freudenburg et al., 2024).

Trust also has a temporal dimension. Initial confidence may be based on a clinical recommendation or a technological promise. Still, durable trust requires reliable maintenance, data protection, honest communication about uncertainty, and clear plans for discontinuation or post-trial responsibility. This is particularly relevant in implanted or long-term systems, where trust becomes part of an ongoing relationship between the user and the institutions responsible for the technology (Baker et al., 2023; Starke & Ienca, 2025).

The Relationship Between Ethics and Trust in the Context of BCI Technology

In the context of BCI technology, ethics and trust are best understood as mutually reinforcing rather than conceptually separate. Ethical principles define the normative conditions under which the development and use of BCIs are legitimate. At the same time, trust reflects whether those conditions are recognised and socially validated by users and the broader public. Ethics provides the standards, whereas trust reflects the extent to which technological practice is perceived as meeting them (Yuste, 2023; Starke & Ienca, 2025).

This relationship is especially evident in data governance. Ethical norms require that neurodata be handled with exceptional care because of their sensitivity and the consequences of misuse. From the standpoint of trust, however, it is not enough for institutions merely to claim that such data are protected. Users and the public must have reason to believe that these protections are real, enforceable, and transparent. Trust, therefore, depends on the practical implementation of ethical commitments through clear consent procedures, security safeguards, purpose limitation, and accountable oversight (Yuste, 2023; Poo, 2024).

A similar connection appears in questions of autonomy and user control. Ethical discussion emphasises that users should remain meaningful agents in relation to BCI outputs, especially in systems that mediate communication or action through machine learning. Trust becomes possible when users can understand the technology's role, influence its operation, and contest outputs that do not reflect their intentions. When systems become opaque or overly predictive, trust is weakened because the ethical basis for

relying on the technology is also weakened (van Stuijvenberg et al., 2024; Soldado-Magraner et al., 2024).

The relationship between ethics and trust is also visible at the societal level. Public trust in BCI technology depends not only on what individual devices achieve, but also on whether the broader trajectory of innovation appears aligned with the public interest. Commercial overstatement, unequal access, exploitative business models, or failures of post-trial responsibility can therefore undermine trust far beyond a single product or institution. In this sense, trust should be understood as one of the most visible social expressions of ethically responsible innovation (Pfotenhauer et al., 2021; Boonstra, 2025; Starke & Ienca, 2025).

METHODOLOGY

Research Approach

This article adopts a theoretical review design based on a structured narrative approach to academic literature. Its purpose is not to test causal hypotheses empirically or to compare effect sizes across studies, but to synthesise, organise, and interpret recent scholarship to clarify the conceptual structure of ethical debates surrounding brain-computer interfaces. Such an approach is particularly appropriate when the field under examination is interdisciplinary, normatively complex, and characterised more by conceptual fragmentation than by methodological uniformity. In the case of BCI ethics, relevant discussions are dispersed across neuroethics, neuroscience, biomedical engineering, medical law, philosophy of technology, and science and technology governance (Snyder, 2019; Baethge et al., 2019).

The choice of a theoretical review is consistent with the article's aims. The study seeks to identify the ethical domains that dominate recent BCI scholarship, examine how these concerns shift across different contexts of use, and determine which governance principles emerge as normatively significant. These questions require a methodology capable of comparing arguments, tracing recurring concepts, and identifying areas of convergence and divergence across publications. A structured narrative review is therefore well-suited to the present analysis because it allows for conceptual clarification and interpretive synthesis rather than empirical generalisation (Snyder, 2019).

Although the article does not follow the strict protocol of a systematic review, the review process was conducted in a transparent and methodologically consistent manner. The study follows a clearly defined logic of source selection, inclusion criteria, thematic organisation, and analytical interpretation. In addition to the analytical literature on BCI ethics and trust, methodological sources were used to support the review design and justify the chosen approach (Baethge et al., 2019; Snyder, 2019).

Data Collection Methods

The analytical corpus for this article was assembled through a targeted search of recent academic literature on brain-computer interfaces and related ethical issues. Literature identification focused on major scholarly databases and interdisciplinary search platforms, including Scopus, Web of Science, PubMed, and Google Scholar. The search process combined general and topic-specific terms such as “brain-computer interface,” “BCI,” “BCI ethics,” “neuroethics,” “neurodata privacy,” “autonomy,” “agency,” “trust,” “governance,” “commercialisation,” and “cognitive enhancement.” These terms were

selected to capture both core BCI scholarship and adjacent discussions directly relevant to the ethical use of BCI systems.

To improve transparency, the selection process was organised as a structured and sequential procedure. An exploratory search was initially conducted to identify the main ethical themes associated with BCIs. Based on this preliminary mapping, the search terms were refined to focus on the themes most relevant to the research questions, including neurodata privacy, autonomy, agency, user control, trust, governance, commercialisation, and cognitive enhancement. Potentially relevant publications were then screened by title and abstract, followed by full-text reading when they appeared to make a substantive ethical or conceptual contribution. The final corpus was selected purposively based on relevance, thematic coverage, publication date, peer-reviewed status, and traceability via DOI information. The main stages of the literature selection procedure are summarised in Table 1.

Table 1. Overview of the Literature Selection Procedure

Stage	Procedure	Purpose
Identification	Search in Scopus, Web of Science, PubMed, and Google Scholar	To locate recent interdisciplinary literature on BCI ethics and trust
Screening	Title, abstract, and full-text relevance assessment	To exclude purely technical or only indirectly relevant publications
Eligibility	Peer-reviewed, English-language, DOI-traceable sources from 2018–2025	To ensure academic quality and traceability
Inclusion	Final purposive corpus of 15 publications	To enable focused conceptual synthesis

Because the aim of the article is theoretical synthesis rather than aggregating systematic evidence, the corpus was not intended to be statistically representative of the entire BCI ethics literature. The relatively small final corpus allows for closer conceptual reading, but also limits the breadth of coverage. For this reason, the findings should be interpreted as an analytically focused reconstruction of dominant ethical patterns in recent scholarship rather than as an exhaustive mapping of the field.

The search was limited to publications from 2018 to 2025. This time frame was chosen because it reflects the contemporary phase of BCI development, characterised by growing clinical translation, increased public and commercial interest, and the rising prominence of issues such as artificial intelligence, speech decoding, neurodata governance, and enhancement-oriented applications. Focusing on this period helped ensure that the analysis remained centred on current conceptual developments rather than earlier debates framed under substantially different technological conditions.

Source selection was guided by four inclusion criteria. Publications had to address BCIs directly or discuss a neurotechnological issue with clear relevance to BCI ethics; make a substantive conceptual, normative, or analytical contribution to themes such as privacy, autonomy, identity, trust, justice, commercialisation, or governance; be published in English in a peer-reviewed academic outlet; and include a verifiable DOI to support traceability and bibliographic accuracy. Publications were excluded if they were purely technical without ethical relevance, lacked sufficient academic grounding, fell outside the selected time frame, or addressed neurotechnology only indirectly. These criteria applied to the analytical corpus, while methodological and contextual sources were used only to support the review design and to illustrate its practical application.

The final analytical corpus comprised 15 peer-reviewed publications on BCI ethics and trust. After the initial search, potentially relevant sources were screened for relevance based on title, abstract, and full text. Backward citation tracing was also selectively used to identify influential recent publications that were repeatedly cited across the literature and clearly aligned with the study's thematic scope. In addition, two methodological sources and several supplementary contextual sources were included to support the review design, clarify selected thematic areas, and illustrate recent BCI developments and governance frameworks. These supplementary sources were not treated as part of the 15-publication analytical corpus. The analytical corpus, therefore, refers only to peer-reviewed publications that directly contributed to the thematic synthesis of BCI ethics and trust, while methodological, illustrative, and policy-oriented sources were used in a supporting role. The final selection was therefore purposive and conceptually focused rather than exhaustive in the systematic-review sense, consistent with the aims of a theoretical article that prioritises depth of interpretation over bibliometric completeness.

Data Analysis Methods

The analysis followed an interpretive thematic procedure designed to identify recurring ethical patterns in recent BCI scholarship. Once the final corpus had been established, each publication was read in full with attention to its central ethical claims, conceptual vocabulary, argumentative structure, and contextual assumptions. The first stage of analysis involved identifying both explicit and implicit ethical concerns, including privacy, informed consent, autonomy, agency, identity, fairness, commercialisation, accountability, and regulation. At this stage, the aim was not to impose a predetermined framework, but to remain open to the diversity of concepts and emphases present across the literature.

In the second stage, the identified concerns were grouped into broader thematic categories through comparative reading. This made it possible to distinguish between recurring ethical domains and more context-specific issues. The analysis showed that recent BCI scholarship clusters most consistently around several interconnected themes: neurodata privacy and data governance; autonomy, informed consent, and meaningful user control; identity, agency, and psychological continuity; justice, access, and commercialisation; and trust and governance as cross-cutting concerns. Organising the material in this way allowed the study to move beyond a source-by-source description toward a more integrated understanding of the field's conceptual structure.

In the third stage, these thematic findings were interpreted in relation to the article’s three research questions. This part of the analysis followed a funnel-shaped logic, moving from general mapping toward more specific normative synthesis. The study first identified the ethical domains that appear most central in current BCI scholarship. It then examined how these concerns change across therapeutic, enhancement-oriented, and commercial contexts of use. Finally, it analysed which governance principles emerge most consistently as conditions for ethically acceptable development and implementation. In this sense, the research questions did not merely organise the presentation of findings, but also guided the transition from description to interpretation.

Because the article is theoretical in nature, the purpose of the analysis was not to quantify the frequency of concepts or to rank studies by methodological strength, but to clarify relationships among arguments and to reconstruct the normative architecture of the field. Particular attention was paid to points of convergence, unresolved tensions, and shifts in emphasis across application contexts. The contribution of the analysis, therefore, lies in showing that contemporary BCI ethics can be understood as a coherent,

though internally differentiated, field in which privacy, autonomy, identity, trust, justice, and governance are best interpreted as interdependent rather than isolated concerns (Snyder, 2019).

RESULTS

As this article is based on a theoretical review rather than an empirical dataset, the results presented in this section should be understood as findings of conceptual synthesis. The section highlights the main ethical and trust-related patterns identified in recent BCI scholarship across different contexts of use.

Presentation of the Key Findings

The review shows that contemporary scholarship no longer treats brain-computer interfaces solely as technical or clinical tools, but increasingly as technologies with biomedical, social, legal, and political significance. BCIs are understood as systems that mediate relations between users, clinicians, developers, institutions, and markets, which means that their ethical evaluation depends not only on performance but also on governance, data practices, and the distribution of control and responsibility (Cabrera & Weber, 2023; Soldado-Magraner et al., 2024).

Recent BCI scholarship also converges around a relatively stable set of ethical domains. Privacy, autonomy, agency, identity, justice, governance, and trust repeatedly emerge as central concerns. These are best understood as interdependent dimensions of a broader ethical field rather than as isolated issues (Yuste, 2023; van Stuijvenberg et al., 2024).

The review also shows that ethical concerns vary across application contexts. Therapeutic uses are mainly associated with vulnerability, informed consent, long-term support, and the preservation of agency. In contrast, enhancement-oriented and commercial applications place greater emphasis on social pressure, fairness, weak evidentiary standards, and the commodification of neural data. Trust, in turn, emerges as a central condition of social legitimacy, since it depends not only on technical functionality, but also on transparency, accountability, and reliable governance structures (Gordon & Seth, 2024; Boonstra, 2025; Starke & Ienca, 2025).

A further finding is that the literature contains important tensions, as well as areas of agreement. These tensions concern the relationship between therapeutic promise and ethical caution, the relationship between innovation and commercialisation, and the status of trust itself. The same features that make BCIs valuable in clinical settings, such as direct access to neural signals and personalised adaptive decoding, also intensify concerns about privacy, dependency, and user control. Commercial investment may accelerate development and translation, but it may also encourage premature claims, proprietary data practices, and unequal access. Trust is necessary for adoption, but it can become ethically problematic when produced through hype, branding, or unrealistic expectations rather than through transparent evidence, accountable governance, and meaningful user involvement. These tensions indicate that BCI ethics should not be reduced to a checklist of principles; it requires critical evaluation of how principles interact in concrete contexts of use.

Findings Regarding the Ethical Aspects of the Use of BCI Technology

The review identifies neurodata privacy as one of the most fundamental concerns in contemporary BCI ethics. Neural data are widely treated as ethically distinct from ordinary digital traces because they are closely linked to cognition, intention, affective states, and communication. For this reason, privacy

concerns extend beyond unauthorised access to include storage, secondary use, commercialisation, cross-platform sharing, and insufficient regulation of brain-derived information. The literature therefore suggests that mental privacy and neurodata governance have become central rather than secondary issues in the ethical evaluation of BCI systems (Yuste, 2023; Poo, 2024).

Another major finding concerns autonomy, informed consent, and meaningful user control. Across the reviewed literature, autonomy is not understood as a one-time act of authorisation, but as an ongoing condition that must be preserved throughout the lifecycle of BCI use. This is especially important in systems integrated into communication, mobility, or long-term care, where future consequences are difficult to anticipate. In AI-driven and speech-related BCIs, these concerns become even more pronounced because predictive models may influence user expression, making meaningful control over outputs a central ethical issue (Kögel et al., 2020; van Stuijvenberg et al., 2024).

The literature also confirms the importance of identity, agency, and psychological continuity. BCIs may be experienced as empowering when they restore communication or control. Still, they may also be disruptive when outputs become uncertain, when agency is partly redistributed to algorithmic systems, or when device removal interrupts an established mode of action. Ethical reflection on BCIs must therefore address not only regulation and safety, but also the user's lived experience of authorship, embodiment, and selfhood (Gilbert et al., 2023; Kögel et al., 2020).

Further findings concern justice, access, commercialisation, and governance. The literature repeatedly warns that BCI development may deepen inequalities if access remains limited, if commercial incentives override therapeutic priorities, or if enhancement-oriented uses create new pressures for optimisation. At the same time, responsibility for BCI systems is distributed among researchers, clinicians, engineers, companies, and regulators, making accountability difficult to trace. Ethically acceptable BCI development, therefore, requires governance mechanisms capable of protecting neurodata, sustaining meaningful consent, clarifying responsibility, and aligning innovation with public values (Boonstra, 2025; Pfotenhauer et al., 2021; Soldado-Magraner et al., 2024; Haag et al., 2025; Zhang et al., 2024).

Findings Regarding Trust in BCI Technology

The review indicates that trust in BCI technology is treated in the literature as a multidimensional and relational construct. Rather than being reduced to optimism or a positive attitude, trust is associated with perceptions of competence, reliability, intelligibility, and accountability. In the context of BCIs, this is particularly important because the technology operates in areas closely tied to communication, vulnerability, bodily integrity, and mental privacy (Starke & Ienca, 2025; Cabrera & Weber, 2023).

The literature further shows that trust operates on several levels. At the individual level, it concerns whether users can rely on the system without fearing loss of control, unpredictable outputs, or data misuse. At the clinical level, it concerns relationships between users and professionals. At the institutional and societal level, it concerns whether companies, regulators, and research organisations behave transparently and responsibly. These levels are interconnected, since weak governance or exaggerated public claims can undermine user trust even when a device performs well in controlled settings (Pfotenhauer et al., 2021; Boonstra, 2025; Starke & Ienca, 2025).

A particularly consistent finding is that user control is one of the strongest determinants of trust in advanced BCI systems. This is especially evident in literature on AI-driven speech BCIs, where trust

depends not only on whether the system can generate useful outputs, but also on whether the user can meaningfully shape what is said and how closely the final output reflects the user's intentions. When predictive models play a substantial role, trust becomes closely tied to transparency and contestability (van Stuijvenberg et al., 2024).

The review also shows that trust can be undermined not only by technical failure, but also by hype, overpromising, and commercialisation without adequate safeguards. Premature market expansion may create misplaced trust when BCI technologies are presented as more mature or transformative than current evidence supports. Overall, the literature suggests that trust is closely linked to broader questions of data governance, user control, and institutional responsibility (Yuste, 2023; Boonstra, 2025; Starke & Ienca, 2025).

DISCUSSION

The findings can be synthesised into a context-sensitive legitimacy framework for BCI ethics. According to this framework, the ethical acceptability of a BCI system depends on four interacting dimensions: the purpose of use, the vulnerability and expectations of users, the degree of algorithmic and institutional mediation, and the governance of neural data and responsibility. This framework helps explain why the same ethical principle may have different practical implications across contexts. For example, autonomy in therapeutic use may require supported decision-making, post-trial responsibility, and continuity of care. In contrast, autonomy in consumer settings may require protection against manipulative marketing, opaque data practices, and weak consent mechanisms. Similarly, trust in clinical contexts depends heavily on professional responsibility and long-term support, while trust in commercial contexts depends on evidence standards, transparency, and limits on data exploitation.

Discussion of RQ1

The first research question asked which ethical domains are most central in contemporary academic discussions of brain-computer interfaces. Based on the reviewed literature, the answer is clear: current BCI ethics is structured around recurring concerns involving privacy and neurodata governance, autonomy and meaningful user control, identity and agency, justice and access, and broader questions of governance and accountability. These domains do not appear as isolated themes but as conceptually connected and practically interdependent, indicating that BCI ethics has developed into a more coherent normative field.

From a theoretical perspective, this convergence is not surprising. BCIs are ethically distinctive because they engage with signals closely linked to intention, cognition, affect, and communication. This gives them a different normative profile from many other digital technologies. Privacy, therefore, becomes especially important, since neural data are widely regarded as highly sensitive. In contrast, autonomy becomes more complex because the key issue is not only whether a user consents to the system, but whether the user remains a meaningful agent once computational processes partly mediate interpretation and output generation.

The literature also shows that identity and agency should be treated as central rather than peripheral concerns. Questions of risk, safety, and feasibility often dominated earlier debates. However, more recent scholarship suggests that these categories alone do not fully capture how BCIs may reshape self-experience, communicative authorship, and dependence on technology. In this sense, the field has

moved beyond a narrowly biomedical perspective toward a broader neuroethical framework attentive to lived experience and self-determination.

Justice likewise emerges as a foundational ethical domain rather than a supplementary policy concern. Questions of access, commercialisation, expectation management, and market pressure are no longer external to BCI ethics, but part of its core structure. The literature repeatedly warns that BCI development may reproduce inequalities, privilege commercially attractive applications over urgent therapeutic needs, and normalise premature translation into consumer markets. From this perspective, the answer to RQ1 is not merely a list of ethical categories, but the broader conclusion that contemporary BCI ethics is organised around an integrated framework in which privacy, autonomy, identity, justice, and governance function as mutually reinforcing dimensions of legitimacy.

An important implication of this finding is that trust should not be placed outside this ethical architecture. The literature suggests that trust operates as a cross-cutting indicator of whether these domains are being adequately addressed. Where privacy is weak, governance is opaque, and user control is limited, trust is likely to remain unstable. Conversely, when ethical safeguards are embedded in design and oversight, trust becomes more justified and durable. This shows that the central ethical domains identified in response to RQ1 are not only analytically important but also socially consequential because they shape whether BCI technologies are regarded as worthy of reliance in clinical and public life.

Discussion of RQ2

The second research question asked how ethical concerns related to BCI technology change across therapeutic, enhancement-oriented, and commercial contexts of use. The review shows that the ethical significance of BCI technology cannot be assessed in a context-neutral way. The same system may carry very different ethical weight depending on whether it is designed to restore lost communication in a vulnerable patient, enhance performance in healthy users, or generate value in a commercial market. Ethical evaluation must therefore remain sensitive to differences in need, vulnerability, benefit, risk tolerance, and institutional power.

In therapeutic contexts, ethical discussion is primarily anchored in need, dependency, and vulnerability. BCIs are often justified by their potential to restore communication, mobility, or functional participation for individuals with severe neurological impairment. At the same time, the literature makes clear that therapeutic urgency does not remove ethical difficulty. Rather, it increases the importance of robust consent, realistic expectation management, long-term device support, post-trial obligations, and attention to the user's ongoing experience of agency and dependence. The ethical challenge lies not only in whether the device works, but also in whether the surrounding system of care remains responsive to vulnerable users.

The example of speech-related BCIs illustrates this point clearly. Recent speech neuroprostheses demonstrate that BCIs may support text generation, synthetic voice output, and avatar-mediated communication for people with severe paralysis or speech loss (Metzger et al., 2023; Wairagkar et al., 2025). In therapeutic contexts, the ability to restore communication may justify substantial technical and clinical burdens. However, the same application also raises unusually sensitive questions about communicative authorship. If an AI-assisted BCI predicts words, tone, or expressive features, the ethical question is not only whether the output is accurate, but whether the user can recognise it as their own. Misclassification in a conventional assistive device may be inconvenient; misclassification in a speech-

BCI may affect self-expression, social identity, and interpersonal relationships. This example shows why therapeutic benefit does not eliminate ethical risk, but changes its structure.

In enhancement contexts, the structure of ethical evaluation changes significantly. Because the benefits are usually less urgent and more contestable than in therapeutic use, concerns about authenticity, coercive social norms, competitive inequality, and the meaning of human agency become more prominent. Enhancement-oriented BCIs, therefore, reveal tensions between innovation and fairness, especially when social or cultural pressure encourages adoption outside medical necessity.

Commercial and consumer contexts intensify these tensions further. Here, the literature highlights weak evidence, persuasive marketing, hype, data extraction, and regulatory asymmetry. The ethical concern is no longer centred mainly on a clinical relationship, but on market structures in which users may simultaneously function as customers, experimental subjects, and sources of valuable neural data. In such settings, consent may be weakened by exaggerated claims, while trust may be shaped more by branding than by transparent evidence and accountable governance.

Taken together, the discussion of RQ2 shows that the ethical profile of BCIs changes in a patterned way across contexts. Therapeutic use foregrounds vulnerability and continuity of care, enhancement foregrounds authenticity and fairness, and commercial use foregrounds governance deficits and the risk of misplaced trust. The broader implication is that BCI ethics should resist one-size-fits-all evaluation. Ethical permissibility depends not only on what the technology does, but also on why it is used, who bears its risks, who benefits from it, and which institutions shape its use.

Discussion of RQ3

The third research question asked which governance principles emerge from recent scholarship as necessary for the ethically acceptable development and implementation of BCI technology. The reviewed literature points to a relatively consistent answer: BCI governance must be proactive, context-sensitive, transparent, and human-centred. It cannot be reduced to formal compliance or post hoc correction after harm occurs. Because BCIs operate in domains tied to mental privacy, communicative authorship, bodily integrity, and long-term dependency, governance must extend from early design and research through clinical translation, long-term use, and eventual discontinuation. Recent work on the ethical review of clinical BCI research reinforces this conclusion by emphasising the need for more explicit governance standards in medical and healthcare applications of BCI technology (Wang et al., 2024).

This conclusion is consistent with recent developments in international governance. The Recommendation of the Council on Responsible Innovation in Neurotechnology, issued by the Organisation for Economic Co-operation and Development (OECD, 2019), emphasises safety assessment, inclusivity, societal deliberation, oversight capacity, protection of personal brain data, stewardship, trust, and anticipation of unintended use or misuse. More recently, the Recommendation on the Ethics of Neurotechnology, adopted by the United Nations Educational, Scientific and Cultural Organization (UNESCO, 2025), has reinforced the need to treat neurotechnology governance as a human-rights-oriented, human-centred, multi-stakeholder, and lifecycle-based responsibility. These policy developments support the article's argument that BCI governance should not be limited to research ethics approval or device safety. It must also address data governance, commercial responsibility, social justice, public deliberation, long-term accountability, and the prevention of misuse.

A recurring theme in the literature is that consent must be understood as ongoing rather than as a one-time event. In BCI settings, risks, uses, and personal meanings may change over time, especially as users encounter issues related to maintenance, data use, changing capabilities, or dependence on the device. For this reason, recent scholarship supports dynamic and revisitable consent, combined with communication practices adapted to users' needs and vulnerabilities.

Another prominent principle concerns neurodata governance, including privacy-by-design. Neural data are increasingly treated as especially sensitive and therefore require strict safeguards for collection, storage, access, purpose limitation, and secondary use. This is particularly important in light of commercial development and the broader data economy, where neural data may be aggregated, monetised, or repurposed beyond the user's reasonable expectations. Privacy is therefore not simply a legal addition to BCI development, but a constitutive condition of ethical legitimacy and justified trust.

The literature also emphasises the importance of meaningful user control and transparent accountability. In advanced BCIs, especially speech-generating and AI-assisted systems, ethical governance requires more than technical accuracy. Users must retain a meaningful role in shaping outputs, correcting errors, and contesting system behaviour. At the same time, responsibility for harm or malfunction must not become so widely distributed across developers, clinicians, institutions, and algorithms that accountability disappears. Governance frameworks are therefore expected to define roles and responsibilities clearly.

Questions of fairness in access, benefit distribution, and expectation management are equally central. Ethical governance cannot focus only on individual protection while ignoring structural inequalities. If BCI innovation remains available only to privileged groups, if public narratives systematically overstate likely benefits, or if commercial incentives drive deployment faster than evidence and oversight allow, governance has failed in a broader social sense. Responsible BCI development, therefore, requires distributive justice, realistic communication, and protection against exploitative or premature commercialisation.

Taken together, the discussion of RQ3 suggests that governance in the BCI field should not be seen as an obstacle to innovation, but as the framework through which innovation becomes ethically defensible and socially sustainable. Ongoing consent, privacy-by-design, meaningful user control, transparent accountability, and justice-oriented oversight emerge as the most consistent principles across the reviewed scholarship.

CONCLUSION

Summary of the Key Findings

This article provided a theoretically grounded analysis of the ethical challenges associated with brain-computer interfaces in contemporary academic scholarship. The review showed that BCI ethics has developed into a coherent field structured around several recurring domains: neurodata privacy, autonomy and meaningful user control, identity and agency, justice and access, and governance. These domains are interconnected and together shape the ethical legitimacy of BCI development and use.

The findings also indicate that ethical concerns vary across contexts. In therapeutic settings, they are mainly related to vulnerability, restoration of lost function, continuity of care, and preservation of

agency. In enhancement contexts, the literature places greater emphasis on authenticity, fairness, and social pressure. In commercial settings, the dominant concerns include hype, weak evidentiary standards, data commodification, and power asymmetries between users and market actors. This suggests that BCI ethics requires a context-sensitive rather than universal approach.

The central argument of this article is that BCI ethics should be understood as a question of context-sensitive legitimacy. BCIs cannot be evaluated solely by asking whether they are technically effective or produce therapeutic benefit. They must also be evaluated according to whether users retain meaningful control, whether neural data are protected, whether responsibility is clearly allocated, whether expectations are realistic, and whether access and benefits are distributed fairly. Trust is therefore not merely a desirable outcome of BCI development, but a test of whether ethical safeguards are sufficiently embedded in technological design, institutional practice, and governance. In this sense, ethics is not external to BCI innovation; it is one of the conditions under which such innovation becomes socially defensible.

Limitations of the Study

This study has several limitations. It is based on a theoretical review rather than on primary empirical research. Its conclusions, therefore, reflect an interpretive synthesis of existing scholarship rather than direct evidence from BCI users, clinicians, developers, or regulators. The article is consequently better suited to conceptual clarification than to empirical generalisation.

The review is also based on a purposive corpus of 15 peer-reviewed publications rather than a fully systematic or bibliometric review. This corpus was selected to support focused conceptual analysis across privacy, autonomy, identity, trust, justice, commercialisation, and governance. However, the limited size of the corpus means that some relevant contributions may not have been included. The findings should therefore be understood as a structured theoretical synthesis of dominant ethical patterns, not as an exhaustive representation of the entire field.

The practical examples discussed in the article are illustrative rather than empirical case studies. They are used to show how ethical concerns become concrete in real or emerging BCI contexts, but they do not replace qualitative research with users or clinical stakeholders. Future research should therefore test and refine the proposed context-sensitive legitimacy framework through empirical studies.

BCI ethics is also a rapidly evolving field. Advances in AI-driven speech decoding, closed-loop neurotechnology, implantable systems, and consumer-facing neurotechnology may shift the relative importance of ethical concerns. The conclusions of this article should therefore be treated as a contribution to an ongoing debate rather than as a final framework.

Suggestions for Future Research

Future research should develop in several directions. More empirical work is needed on how users, families, clinicians, and developers experience ethical challenges in real BCI contexts, especially over time. Qualitative and mixed-method studies could make an important contribution by examining how ethical issues evolve across the full lifecycle of BCI use, including adoption, long-term integration, maintenance, and discontinuation.

Further research should also examine the ethical implications of artificial intelligence in BCI systems, particularly regarding explainability, user control, responsibility, and trust. This is especially important in speech-generating and adaptive BCI systems, where predictive models may increasingly shape communication and action.

Finally, more attention should be given to justice and global inclusion. Future scholarship should examine who benefits from BCI technologies, whose needs shape research priorities, and whether governance frameworks are responsive to diverse legal, cultural, and socioeconomic contexts. A more comparative and justice-oriented research agenda would strengthen both the academic depth and the practical relevance of future work in this field.

ACKNOWLEDGEMENT

The authors acknowledge the use of digital tools for grammatical refinement and structural organisation during the preparation of this manuscript. These tools supported clarity and consistency in presentation. However, the arguments, interpretations, conclusions, and any remaining errors are solely the responsibility of the authors.

REFERENCES

Baethge, C., Goldbeck-Wood, S., & Mertens, S. (2019). SANRA—a scale for the quality assessment of narrative review articles. *Research Integrity and Peer Review*, 4, Article 5. https://doi.org/10.1186/s41073-019-0064-8

Baker, S., Fenstermacher, E., Davis, R. A., Kern, D. S., Thompson, J. A., Felsen, G., & Baumgartner, A. J. (2023). Ethical considerations in closed-loop deep-brain stimulation. *Deep Brain Stimulation*, 3, 8–15. https://doi.org/10.1016/j.jdbs.2023.11.001

Boonstra, J. T. (2025). Ethical imperatives in the commercialisation of brain-computer interfaces. *IBRO Neuroscience Reports*, 19, 718–724. https://doi.org/10.1016/j.ibneur.2025.10.004

Cabrera, L. Y., & Weber, D. J. (2023). Rethinking the ethical priorities for brain–computer interfaces. *Nature Electronics*, 6(2), 99–101. https://doi.org/10.1038/s41928-023-00928-w

Freudenburg, Z. F., Berezutskaya, J., & Herbert, C. H. (2024). Editorial: The ethics of speech ownership in the context of neural control of augmented assistive communication. *Frontiers in Human Neuroscience*, 18, 1468938. https://doi.org/10.3389/fnhum.2024.1468938

Gilbert, F., Ienca, M., & Cook, M. (2023). How I became myself after merging with a computer: Does human-machine symbiosis raise human rights issues? *Brain Stimulation*, 16(3), 783–789. https://doi.org/10.1016/j.brs.2023.04.016

Gordon, E. C., & Seth, A. K. (2024). Ethical considerations for the use of brain–computer interfaces for cognitive enhancement. *PLoS Biology*, 22(10), e3002899. https://doi.org/10.1371/journal.pbio.3002899

Haag, L., Starke, G., Ploner, M., & Ienca, M. (2025). Ethical gaps in closed-loop neurotechnology: A scoping review. *npj Digital Medicine*, 8, Article 510. https://doi.org/10.1038/s41746-025-01908-4

Kögel, J., Jox, R. J., & Friedrich, O. (2020). What is it like to use a BCI? Insights from an interview study with brain-computer interface users. *BMC Medical Ethics*, 21, Article 2. https://doi.org/10.1186/s12910-019-0442-2

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(7976), 1037–1046. https://doi.org/10.1038/s41586-023-06443-4

Organisation for Economic Co-operation and Development. (2019). *Recommendation of the Council on Responsible Innovation in Neurotechnology* (OECD/LEGAL/0457). OECD Legal Instruments. https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0457

Pfotenhauer, S. M., Frahm, N., Winickoff, D., Benrimoh, D., Illes, J., & Marchant, G. (2021). Mobilizing the private sector for responsible innovation in neurotechnology. *Nature Biotechnology*, 39(6), 661–664. https://doi.org/10.1038/s41587-021-00947-y

Poo, M.-M. (2024). China's new ethical guidelines for the use of brain–computer interfaces. *National Science Review*, 11(4), nwae154. https://doi.org/10.1093/nsr/nwae154

Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. https://doi.org/10.1016/j.jbusres.2019.07.039

Soldado-Magraner, J., Antonietti, A., French, J., Higgins, N., Young, M. J., Larrivee, D., & Monteleone, R. (2024). Applying the IEEE BRAIN neuroethics framework to intra-cortical brain-computer interfaces. *Journal of Neural Engineering*, 21(2), 022001. https://doi.org/10.1088/1741-2552/ad3852

Specker Sullivan, L., & Illes, J. (2018). Ethics in published brain–computer interface research. *Journal of Neural Engineering*, 15(1), 013001. https://doi.org/10.1088/1741-2552/aa8e05

Starke, G., & Ienca, M. (2025). AI, neurotechnology and society — a question of trust. *Nature Reviews Neurology*, 21, 407–408. https://doi.org/10.1038/s41582-025-01114-6

United Nations Educational, Scientific and Cultural Organization. (2025). *Recommendation on the ethics of neurotechnology*. UNESCO. https://www.unesco.org/en/legal-affairs/recommendation-ethics-neurotechnology

van Stuijvenberg, O. C., Samlal, D. P. S., Vansteensel, M. J., Broekman, M. L. D., & Jongsma, K. R. (2024). The ethical significance of user-control in AI-driven speech-BCIs: A narrative review. *Frontiers in Human Neuroscience*, 18, Article 1420334. https://doi.org/10.3389/fnhum.2024.1420334

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. https://doi.org/10.1038/s41586-025-09127-3

Wang, X.-Q., Sun, H.-Q., Si, J.-Y., Lin, Z.-Y., Zhai, X.-M., & Lu, L. (2024). Challenges and suggestions of ethical review on clinical research involving brain-computer interfaces. *Chinese Medical Sciences Journal*, 39(2), 131–139. https://doi.org/10.24920/004377

Wexler, A., & Thibault, R. (2019). Mind-reading or misleading? Assessing direct-to-consumer electroencephalography (EEG) devices marketed for wellness and their ethical and regulatory implications. *Journal of Cognitive Enhancement*, 3(1), 131–137. https://doi.org/10.1007/s41465-018-0091-2

Yuste, R. (2023). Advocating for neurodata privacy and neurotechnology regulation. *Nature Protocols*, 18(10), 2869–2875. https://doi.org/10.1038/s41596-023-00873-0

Zhang, Z., Chen, Y., Zhao, X., Fan, W., Peng, D., Li, T., Zhao, L., & Fu, Y. (2024). A review of ethical considerations for the medical applications of brain-computer interfaces. *Cognitive Neurodynamics*, 18, 3603–3614. https://doi.org/10.1007/s11571-024-10144-7