

Neural Interfaces and Human Augmentation: Technological Opportunities, Ethical Challenges and Social Implications

Nicolas Aumonier

Professor, University Grenoble Alpes, France

Abstract

Neural interfaces are transforming the relationship between the human nervous system and digital technology. Brain–computer interfaces, neuroprosthetic devices, closed-loop stimulation systems, wearable neural sensors, and artificial intelligence-assisted decoding platforms can restore communication, movement, sensory perception, and functional independence for people affected by neurological injury or disease. The same technological foundation may also support human augmentation by extending memory, attention, perception, motor control, and interaction with computational systems. These developments create substantial medical and social opportunities, but they also introduce concerns involving mental privacy, cognitive liberty, personal identity, informed consent, cybersecurity, distributive justice, and the boundary between therapy and enhancement.

This review examines neural interfaces through an interdisciplinary framework connecting biomedical engineering, neuroscience, artificial intelligence, bioethics, disability studies, and technology governance. A structured conceptual synthesis was conducted using established scholarly literature published up to 2021. The analysis organized the evidence into five domains: therapeutic restoration, cognitive and sensory augmentation, neural-data governance, human agency, and social accessibility. A transparent priority-scoring framework was used to compare the technological potential and governance sensitivity of major application areas.

The synthesis indicates that assistive communication and motor restoration possess the strongest near-term social justification, whereas workplace monitoring, military enhancement, and consumer cognitive augmentation present particularly serious ethical and governance risks. The paper argues that responsible neural-interface innovation requires privacy-by-design, meaningful consent, user control, long-term safety monitoring, independent oversight, equitable access, and explicit restrictions on coercive or manipulative uses. Neural augmentation should therefore be evaluated not only according to technical performance but also according to its effects on dignity, freedom, social equality, and the conditions under which individuals retain genuine control over their cognitive lives.

Keywords: neural interfaces; brain–computer interfaces; human augmentation; neuroprosthetics; neural data; cognitive liberty; mental privacy; neuroethics.

1. Introduction

Neural interfaces establish a functional connection between the nervous system and an external technological device. Depending on their design, these systems may record neural activity, translate neural signals into commands, stimulate specific neural structures, or create a continuous feedback loop between biological and computational processes. The broad category includes electroencephalography-based brain-computer interfaces, implanted cortical electrodes, deep-brain stimulation systems, peripheral nerve interfaces, cochlear implants, retinal prostheses, wearable neurotechnology, and machine-learning platforms that interpret neural patterns. What once appeared primarily as an experimental approach to neurological rehabilitation is increasingly becoming a wider technological field with implications for health care, communication, education, employment, security, and consumer technology.

The most compelling applications remain therapeutic. A person affected by paralysis may use an interface to select letters, operate a computer, control a robotic limb, or communicate through synthesized speech. Neuroprosthetic systems may help restore aspects of hearing, vision, movement, or somatosensory feedback. Closed-loop stimulation can detect relevant neural states and adjust stimulation in real time, creating possibilities for more responsive management of neurological and psychiatric conditions. These applications demonstrate that neural interfaces can reduce functional limitations and expand individual independence. They also illustrate how human capability is frequently produced through interaction among the body, technology, caregivers, and the surrounding social environment rather than by biological capacity alone.

The distinction between restoration and augmentation, however, is neither fixed nor ethically straightforward. A device developed to restore a lost function may eventually produce capabilities that exceed an assumed biological baseline. Neural systems designed for rehabilitation may be adapted for attention management, accelerated human-computer interaction, expanded sensory perception, or direct control of remote machines. The meaning of augmentation also depends on social context. A technology may be restorative for one user, assistive for another, and performance-enhancing for a third. Ethical analysis must consequently move beyond a simple therapeutic-versus-enhancement division and consider purpose, proportionality, voluntariness, risk, accessibility, and the distribution of benefits.

Neural technologies differ from many other digital systems because they interact with processes connected to intention, emotion, perception, memory, decision-making, and bodily control. Neural data may not provide transparent access to a person's thoughts, but increasingly sophisticated analytical systems can infer patterns related to attention, fatigue, preferences, health conditions, or behavioral responses. If such data are collected in clinical, commercial, employment, educational, or military settings, conventional data-protection mechanisms may not adequately address their sensitivity. Questions concerning who owns neural data, who can interpret them, how long they can be retained, and whether they can be used for secondary purposes are therefore central to responsible innovation.

This paper evaluates the technological opportunities, ethical challenges, and social implications of neural interfaces and human augmentation. It adopts an interdisciplinary perspective because technical performance cannot be separated from human agency, regulatory responsibility, accessibility, and social power. The central position advanced here is that neural interfaces should be assessed according to a dual standard: their capacity to improve human functioning and their ability to preserve

the user's autonomy, dignity, privacy, security, and meaningful control. Technological success without these protections would represent an incomplete and potentially harmful form of innovation.

2. Review of Literature

2.1 Development of Brain–Computer Interfaces and Neuroprosthetic Systems

Early brain–computer interface research established that recorded neural activity could be translated into commands for external devices. Wolpaw et al. (2002) described the principal communication pathway of a brain–computer interface as one that does not depend on the brain's normal peripheral nerves and muscles. This formulation was especially important for individuals with severe motor impairment because it repositioned neural activity as a possible communication channel. Later research expanded the technical field from basic cursor control toward robotic devices, prosthetic limbs, communication systems, and multisensory feedback.

Nicolelis (2003) examined brain–machine interfaces as a developing approach for restoring motor function through the direct translation of neuronal activity into device control. Lebedev and Nicolelis (2017) subsequently reviewed progress in brain–machine interfaces and identified advances in large-scale neural recording, decoding methods, robotic control, and sensory feedback. Their analysis also showed that reliable operation depends on adaptation between the user and the device. The interface is therefore not merely a passive decoder; it is part of a learning relationship in which neural activity, algorithmic classification, and user behavior change over time.

Collinger et al. (2013) demonstrated high-performance control of a robotic arm by a person with tetraplegia using intracortical neural signals. Such studies illustrate the restorative potential of invasive interfaces while also revealing the demanding clinical, engineering, and training requirements involved. Ajiboye et al. (2017) reported a system that combined intracortical recording with functional electrical stimulation to support reaching and grasping in a person with tetraplegia. These developments shifted the field toward closed functional systems in which decoded intention could be translated into movement of the user's own limb.

2.2 Artificial Intelligence, Adaptation, and Augmented Capability

Machine learning plays a central role in identifying useful information within complex and often noisy neural signals. Neural patterns vary across individuals, tasks, environments, and recording sessions. Decoding systems must therefore be calibrated and, in many cases, updated over time. Adaptive algorithms can improve performance, but they also complicate accountability because system behavior may evolve after initial deployment. Errors may arise from signal drift, changes in the user's health, electrode degradation, incomplete training data, environmental interference, or inappropriate algorithmic assumptions.

Cinel et al. (2019) discussed the convergence of brain–computer interfaces and human–machine interaction as a pathway toward cognitive augmentation. Their work identified possible applications involving decision support, attention, perception, and collaborative human–machine systems. Human augmentation in this context does not necessarily mean replacing human judgment. It may involve creating a distributed cognitive system in which human interpretation and machine computation complement one another.

The literature also warns against assuming that technological integration automatically enhances agency. A system may improve task performance while reducing user understanding or increasing dependence on proprietary algorithms.

If the user cannot determine whether an action originated from personal intention, device prediction, automated correction, or external intervention, responsibility becomes difficult to locate. Clausen (2011) argued that neural devices raise important concerns relating to autonomy, informed consent, responsibility, and personal identity. These issues become more pronounced as systems move from simple command translation toward predictive and closed-loop intervention.

2.3 Neuroethics, Rights, and Social Governance

Neuroethical scholarship has emphasized that technologies interacting with neural processes may affect dimensions of personhood that are not adequately protected by ordinary medical-device regulation. Ienca and Andorno (2017) proposed four emerging rights for the neurotechnological era: cognitive liberty, mental privacy, mental integrity, and psychological continuity. Their framework has influenced debates about whether existing human-rights protections require clarification or extension to address neural data and intervention.

Yuste et al. (2017) similarly argued that the development of neurotechnology and artificial intelligence requires attention to privacy, identity, agency, equality, and bias. The authors proposed integrating ethical considerations into technological development rather than treating them as a later regulatory response. Burwell et al. (2017), after reviewing ethical discussions concerning brain–computer interfaces, found recurring concerns involving autonomy, personhood, privacy, responsibility, safety, and justice.

Goering et al. (2021) connected neurotechnology governance with human rights and recommended protections relating to personal identity, free will, privacy, access, and freedom from bias. This approach is especially relevant because neural interfaces may be used outside traditional clinical environments. Consumer devices, educational monitoring tools, workplace systems, and military applications may not receive the same level of ethical review as medical research, even when their consequences for privacy and autonomy are substantial.

3. Objectives of the Study

3.1 Technological Assessment Objectives

The first objective is to examine the principal technological opportunities associated with neural interfaces. The analysis considers therapeutic restoration, communication assistance, motor neuroprosthetics, sensory substitution, closed-loop neuromodulation, cognitive support, and human–machine collaboration. Attention is given to the relationship between technological capability and practical conditions such as reliability, invasiveness, training, maintenance, and long-term support.

A related objective is to distinguish realistic near-term applications from more speculative augmentation claims. Neural interfaces have produced meaningful experimental achievements, but performance under controlled conditions does not automatically translate into safe and accessible everyday use. The paper therefore assesses opportunity in relation to technical maturity and social value rather than novelty alone.

3.2 Ethical and Governance Objectives

The second objective is to identify the main ethical risks created by neural recording, decoding, stimulation, and behavioral adaptation. Particular attention is given to informed consent, mental privacy, agency, identity, cybersecurity, algorithmic bias, reversibility, and the possibility of coercive use. These concerns are examined across clinical and nonclinical environments.

The study also seeks to formulate governance principles capable of supporting innovation without weakening fundamental protections. The objective is not to prohibit human augmentation categorically, but to establish conditions under which applications may be considered ethically defensible, socially beneficial, and subject to meaningful accountability.

3.3 Social Evaluation Objectives

The third objective is to evaluate how neural interfaces may affect equality, disability inclusion, employment, education, and social expectations of human performance. Access to augmentation may become a source of competitive advantage, while the inability or refusal to adopt it could produce new forms of exclusion.

The study accordingly examines whether emerging neural technologies are likely to reduce existing disadvantages or reproduce them through cost, infrastructure, data bias, and unequal institutional power. Social evaluation is treated as part of technological assessment rather than as a secondary consequence.

4. Research Methodology

4.1 Research Design

The study employed a conceptual and evidence-based review design. It did not involve human participants, clinical experimentation, or the collection of primary neural data. The review integrated findings from neuroscience, biomedical engineering, artificial intelligence, neuroethics, law, and disability studies to develop an interdisciplinary evaluation of neural interfaces.

The analytical unit was the application domain rather than an individual device. This approach made it possible to compare technologies according to their functional purpose and governance implications. Seven domains were examined: assistive communication, motor restoration, sensory restoration, therapeutic neuromodulation, cognitive augmentation, workplace neural monitoring, and military or security augmentation.

4.2 Evidence Identification and Selection

Scholarly sources were selected from established literature available through PubMed, IEEE Xplore, Scopus-indexed publications, and multidisciplinary academic databases. Search concepts included “brain–computer interface,” “neural interface,” “neuroprosthesis,” “human augmentation,” “neural data,” “mental privacy,” “cognitive liberty,” “neuroethics,” and “brain–machine interface.” Literature published up to 2021 was considered.

Sources were included when they addressed technological functionality, clinical or assistive applications, ethical concerns, human rights, user experience, data governance, or social consequences. Purely promotional commentary, technically unsupported predictions, and sources without clear

relevance to neural interfaces were excluded. Foundational studies were retained because the field's ethical and technical concepts developed across several decades.

4.3 Analytical Framework

Each application domain was assessed through five dimensions: functional benefit, technical maturity, individual risk, governance sensitivity, and equity concern. Scores were assigned on a five-point ordinal scale. A score of 1 represented a comparatively low level, while a score of 5 represented a comparatively high level. The ratings are interpretive outputs of the literature synthesis and should not be treated as clinical measurements.

Table 1: Analytical Dimensions Used in the Review

Dimension	Analytical question	Score interpretation
Functional benefit	How strongly may the application improve health, communication, independence, or capability?	1 = limited; 5 = transformative
Technical maturity	How close is the application to reliable practical deployment?	1 = speculative; 5 = established
Individual risk	What is the potential for physical, psychological, or autonomy-related harm?	1 = low; 5 = severe
Governance sensitivity	How strongly does the application require specialized consent, privacy, and oversight safeguards?	1 = conventional; 5 = exceptional
Equity concern	Could unequal access or institutional power create exclusion or discrimination?	1 = limited; 5 = substantial

Note: The framework was developed for this conceptual review. Scores summarize comparative judgments based on the reviewed evidence and do not represent experimental outcomes.

4.4 Methodological Limitations

The review synthesizes a rapidly developing and heterogeneous field. Technologies described under the category of neural interfaces differ considerably in invasiveness, purpose, performance, risk, and regulatory status. Direct comparison therefore requires interpretive judgment.

The analytical scores are designed to make that judgment visible rather than to claim mathematical precision. Future systematic reviews could extend the framework through larger databases, multiple independent coders, formal quality appraisal, and consultation with users of assistive and implanted devices.

5. Analysis

5.1 Therapeutic and Assistive Opportunities

Assistive communication represents one of the strongest justifications for neural-interface research. Individuals who retain cognitive capacity but have lost reliable muscular communication may benefit from systems that translate neural activity into letters, words, or device commands. The value of such technology cannot be measured only through decoding accuracy. Independence, social participation, communication speed, fatigue, ease of calibration, caregiver involvement, and the user's ability to employ the device outside a laboratory are equally important.

Motor neuroprostheses may restore functional interaction by translating intended movement into robotic control, computer operation, or electrical stimulation of muscles. Sensory feedback can further improve control by returning information about touch, pressure, or position. This bidirectional architecture offers greater functional integration than one-way command systems, but it also increases system complexity and potential failure points.

Therapeutic neuromodulation presents another important opportunity. Closed-loop devices may monitor neural activity and deliver stimulation only when specified patterns occur. Such responsiveness could improve treatment precision while reducing unnecessary stimulation.

5.2 Augmentation Beyond Clinical Restoration

Neural augmentation may support attention regulation, sensory extension, accelerated interaction with machines, or control of additional digital and robotic systems. Noninvasive interfaces may allow users to interact with virtual environments, robots, or accessibility platforms without conventional physical input. In industrial or high-risk environments, neural and physiological sensing could potentially help identify fatigue or excessive cognitive workload.

The same capacity can be framed very differently depending on institutional power. A voluntary clinical interface intended to restore communication is not ethically equivalent to an employer requiring neural monitoring as a condition of work. Even if both systems record similar signals, the second context contains risks of coercion, surveillance, discriminatory profiling, and secondary data use.

Consumer augmentation is likely to produce additional ambiguity. Marketing may present neural devices as tools for concentration, emotional regulation, learning, or personal optimization before their long-term effectiveness has been adequately established.

5.3 Comparative Opportunity and Governance Assessment


Table 2: Comparative Assessment of Neural-Interface Application Domains

Application domain	Functional benefit	Technical maturity	Individual risk	Governance sensitivity	Equity concern
Assistive communication	5	3	3	4	4
Motor restoration	5	3	4	4	4

Application domain	Functional benefit	Technical maturity	Individual risk	Governance sensitivity	Equity concern
Sensory restoration	4	3	4	4	4
Therapeutic neuromodulation	5	4	4	5	3
Cognitive augmentation	3	2	4	5	5
Workplace neural monitoring	2	2	5	5	5
Military and security augmentation	3	2	5	5	5

The comparison suggests that high functional value does not eliminate ethical sensitivity. Therapeutic neuromodulation, for example, may provide substantial benefit while requiring intensive oversight because stimulation can affect behavior, mood, or cognitive functioning. Assistive communication and motor restoration combine strong social value with significant access concerns, particularly when expensive infrastructure and long-term technical support are required.

Graph 1: Governance Sensitivity of Neural-Interface Applications



Supporting data: Communication = 4; motor restoration = 4; sensory restoration = 4; therapeutic neuromodulation = 5; cognitive augmentation = 5; workplace monitoring = 5; military augmentation = 5.

Interpretation: Applications involving automated neural stimulation, cognitive enhancement, institutional surveillance, or security functions require the most intensive governance. Their sensitivity arises from the combination of neural-data processing, unequal power, potential behavioral influence, and uncertainty about long-term consequences.

Key Finding: Governance requirements should be determined by the nature of neural access and institutional power rather than by whether a product is labeled medical, consumer, educational, or occupational.

6. Challenges

6.1 Technical Reliability, Biological Compatibility, and Long-Term Safety

Neural signals are variable, context-dependent, and difficult to interpret. A decoder trained during one period may perform differently after changes in electrode position, fatigue, medication, emotional state, neural adaptation, or environmental conditions. An error in a consumer interface may be inconvenient, whereas an error in a mobility device, implanted stimulator, vehicle-control system, or industrial application may cause serious harm.

Invasive interfaces introduce additional clinical concerns. Surgical risk, infection, tissue response, electrode degradation, hardware failure, battery replacement, and device removal must be considered throughout the product life cycle. Long-term responsibility is especially important when a manufacturer withdraws support, changes ownership, discontinues a product, or restricts access to proprietary software. Users should not become medically or functionally dependent on technology without enforceable continuity arrangements.

Interoperability is another challenge. Closed proprietary systems can prevent users from moving their data or accessing alternative service providers. Standardized safety reporting, secure software updates, repairability, and clinically supervised transition pathways are necessary if neural devices are expected to function over extended periods.

6.2 Mental Privacy, Cybersecurity, and Neural-Data Control

Neural data are often probabilistic and require interpretation. They should therefore not be described as exact recordings of private thought. Nevertheless, patterns in neural signals may reveal or support inferences about health, attention, emotional response, fatigue, or intended action. When combined with behavioral, biometric, and contextual information, these data can become exceptionally sensitive.

Conventional consent forms may not adequately explain the future analytical uses of neural information. Data initially collected for device calibration could later be used for algorithm development, commercial profiling, insurance decisions, employment assessment, or research. Effective governance requires purpose limitation, data minimization, secure local processing where feasible, controlled access, retention limits, and explicit authorization for secondary use.

Cybersecurity must be incorporated throughout design and maintenance. Unauthorized access could expose private data, interrupt device operation, manipulate output, or alter stimulation parameters.

Security testing should include hardware, software, wireless communication, cloud infrastructure, clinical interfaces, and update mechanisms. Users must also be informed about residual risks in understandable language.

6.3 Consent, Agency, and Personal Identity

Consent for neural interfaces is a continuing process rather than a single signature. Users need information about surgery, device operation, uncertainty, data processing, algorithmic adaptation, maintenance, withdrawal, and possible psychological effects. Consent becomes especially complicated when a person's neurological condition affects communication or decision-making capacity.

Agency concerns arise when systems predict, modify, or complete user actions. If an adaptive decoder corrects a movement before the user recognizes the error, it may improve performance but complicate the experience of authorship. Interfaces should therefore provide adjustable levels of automation, visible operational feedback, accessible interruption controls, and records that allow disputed actions to be examined.

Identity-related concerns should not be reduced to the assumption that every neural device threatens personhood. Many users may experience a prosthesis as enabling a more authentic and independent life. Ethical evaluation should respect these experiences while acknowledging that unwanted behavioral or emotional changes can affect psychological continuity. The decisive issue is whether the user retains meaningful influence over adaptation and continued use.

7. Discussion

7.1 From Technological Performance to Human-Centered Value

The analysis shows that technical accuracy is necessary but insufficient for evaluating neural interfaces. A system may achieve strong laboratory performance while remaining impractical because it requires lengthy calibration, specialist support, restrictive environmental conditions, or frequent maintenance. Human-centered value depends on whether the technology improves the user's daily life under realistic circumstances.

Outcome evaluation should therefore include communication efficiency, fatigue, comfort, reliability, independence, participation, user confidence, and the burden placed on caregivers. Long-term qualitative evidence is particularly important because users may experience benefits and difficulties that cannot be captured through brief technical trials.

The involvement of disabled people and long-term users must extend beyond conventional participant recruitment. Users should contribute to research priorities, interface design, outcome selection, risk evaluation, accessibility standards, and post-market monitoring. This approach can reduce the distance between engineering assumptions and lived experience.

7.2 Therapy, Enhancement, and the Risk of Social Coercion

The boundary between therapy and enhancement cannot independently determine ethical acceptability. Therapeutic interventions can be coercive or poorly governed, while some enhancement applications may be voluntary and comparatively low risk. A more reliable assessment should examine expected benefit, physical risk, reversibility, data sensitivity, voluntariness, institutional power, and social effects.

Competitive environments can convert formally voluntary augmentation into practical coercion. If neural enhancement improves productivity, reaction speed, or attention, workers and students may feel compelled to adopt it to remain competitive. Refusal could be interpreted as lack of commitment or capability. Protection against discrimination for nonuse is therefore as important as access for those who wish to use assistive technology.

Enhancement may also create a self-reinforcing inequality. Individuals with greater resources could obtain early access, specialist training, maintenance, and upgrades, while disadvantaged populations remain excluded. Public policy should prioritize applications that address serious health and accessibility needs before subsidizing elective performance enhancement.

7.3 A Responsible Governance Architecture

A responsible governance model should begin with clear classification of purpose, invasiveness, neural-data sensitivity, and institutional power. High-risk applications require independent ethical review, strong evidence of safety and effectiveness, ongoing monitoring, and accessible complaint mechanisms. Consumer classification should not be used to avoid scrutiny when a device collects sensitive neural information or claims to influence cognition.

Privacy-by-design should include limited collection, local or edge processing where appropriate, encryption, restricted retention, user-accessible records, and meaningful deletion rights. Users should know what data are collected, which inferences are generated, who receives them, and whether they contribute to algorithm training. Consent should be renewed when the purpose or analytical capability changes.

Accountability must extend across the complete system. Device manufacturers, software developers, clinical providers, data processors, research institutions, and deploying organizations each possess distinct responsibilities. Governance should also require continuity plans for maintenance, security updates, technical failure, corporate closure, and safe device discontinuation.

8. Conclusion

Neural interfaces create important possibilities for restoring communication, movement, sensation, and functional independence. Their strongest social justification lies in applications that respond to serious clinical or accessibility needs while giving users greater control over their lives. Progress in neural recording, artificial intelligence, stimulation, and neuroprosthetic engineering may further improve the responsiveness and usability of these systems.

Human augmentation extends the same technological foundation into more contested domains. Cognitive optimization, workplace monitoring, consumer neurotechnology, and military enhancement raise concerns that cannot be resolved through performance testing alone. The ability to record or influence neural activity creates responsibilities concerning mental privacy, agency, identity, informed consent, cybersecurity, and protection from coercion.

The social effects of augmentation will depend substantially on access and institutional power. A device may support autonomy when selected by an informed user yet undermine autonomy when required by an employer, educational institution, insurer, or military authority. Responsible policy must

therefore protect both the right to benefit from assistive neurotechnology and the right to refuse unnecessary neural monitoring or enhancement.

Future development should follow a rights-respecting and user-centered model. This model requires privacy and security by design, independent review, transparent algorithmic operation, meaningful withdrawal options, long-term support, equitable access, and restrictions on manipulative or discriminatory use. Neural interfaces should ultimately be judged by whether they expand human capability without turning cognitive life into an object of surveillance, control, or unequal commercial power.

References

1. Ajiboye, A. B., Willett, F. R., Young, D. R., Memberg, W. D., Murphy, B. A., Miller, J. P., Walter, B. L., Sweet, J. A., Huyen, H. A., Keith, M. W., Peckham, P. H., Simeral, J. D., Donoghue, J. P., Hochberg, L. R., & Kirsch, R. F. (2017). Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: A proof-of-concept demonstration. *The Lancet*, 389(10081), 1821–1830. [https://doi.org/10.1016/S0140-6736\(17\)30601-3](https://doi.org/10.1016/S0140-6736(17)30601-3)
2. Burwell, S., Sample, M., & Racine, E. (2017). Ethical aspects of brain computer interfaces: A scoping review. *BMC Medical Ethics*, 18, Article 60. <https://doi.org/10.1186/s12910-017-0220-y>
3. Clausen, J. (2011). Conceptual and ethical issues with brain–hardware interfaces. *Current Opinion in Psychiatry*, 24(6), 495–501. <https://doi.org/10.1097/YCO.0b013e32834bb8ca>
4. Cinel, C., Valeriani, D., & Poli, R. (2019). Neurotechnologies for human cognitive augmentation: Current state of the art and future prospects. *Frontiers in Human Neuroscience*, 13, Article 13. <https://doi.org/10.3389/fnhum.2019.00013>
5. Collinger, J. L., Wodlinger, B., Downey, J. E., Wang, W., Tyler-Kabara, E. C., Weber, D. J., McMorland, A. J. C., Velliste, M., Boninger, M. L., & Schwartz, A. B. (2013). High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet*, 381(9866), 557–564. [https://doi.org/10.1016/S0140-6736\(12\)61816-9](https://doi.org/10.1016/S0140-6736(12)61816-9)
6. Glannon, W. (2014). Neuromodulation, agency and autonomy. *Brain Topography*, 27(1), 46–54. <https://doi.org/10.1007/s10548-012-0269-3>
7. Goering, S., Klein, E., Specker Sullivan, L., Wexler, A., Agüera y Arcas, B., Bi, G., Carmenta, J. M., Fins, J. J., Friesen, P., Gallant, J., Huggins, J. E., Kellmeyer, P., Marblestone, A., Mitchell, C., Parens, E., Pham, M., Rubel, A., Sadato, N., Teicher, M., ... Yuste, R. (2021). Recommendations for responsible development and application of neurotechnologies. *Neuroethics*, 14, 365–386. <https://doi.org/10.1007/s12152-021-09468-6>
8. Haselager, P. (2013). Did I do that? Brain–computer interfacing and the sense of agency. *Minds and Machines*, 23(3), 405–418. <https://doi.org/10.1007/s11023-012-9298-7>
9. Ienca, M., & Andorno, R. (2017). Towards new human rights in the age of neuroscience and neurotechnology. *Life Sciences, Society and Policy*, 13, Article 5. <https://doi.org/10.1186/s40504-017-0050-1>
10. Kellmeyer, P. (2018). Big brain data: On the responsible use of brain data from clinical and consumer-directed neurotechnological devices. *Neuroethics*, 11, 83–98. <https://doi.org/10.1007/s12152-018-9371-x>

11. Lebedev, M. A., & Nicolelis, M. A. L. (2017). Brain–machine interfaces: From basic science to neuroprostheses and neurorehabilitation. *Physiological Reviews*, 97(2), 767–837. <https://doi.org/10.1152/physrev.00027.2016>
12. Nijboer, F., Clausen, J., Allison, B. Z., & Haselager, P. (2013). The Asilomar survey: Stakeholders' opinions on ethical issues related to brain–computer interfacing. *Neuroethics*, 6, 541–578. <https://doi.org/10.1007/s12152-011-9132-6>
13. Nicolelis, M. A. L. (2003). Brain–machine interfaces to restore motor function and probe neural circuits. *Nature Reviews Neuroscience*, 4(5), 417–422. <https://doi.org/10.1038/nrn1105>
14. Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., & Vaughan, T. M. (2002). Brain–computer interfaces for communication and control. *Clinical Neurophysiology*, 113(6), 767–791. [https://doi.org/10.1016/S1388-2457\(02\)00057-3](https://doi.org/10.1016/S1388-2457(02)00057-3)
15. Yuste, R., Goering, S., Agüera y Arcas, B., Bi, G., Carmena, J. M., Carter, A., Fins, J. J., Friesen, P., Gallant, J., Huggins, J. E., Illes, J., Kellmeyer, P., Klein, E., Marblestone, A., Mitchell, C., Parens, E., Pham, M., Rubel, A., Sadato, N., ... Wolpaw, J. (2017). Four ethical priorities for neurotechnologies and AI. *Nature*, 551(7679), 159–163. <https://doi.org/10.1038/551159a>