

Stanford – Vienna
Transatlantic Technology Law Forum

A joint initiative of
Stanford Law School and the University of Vienna School of Law



TTLF Working Papers

No. 161

Legal and Technical Pathways for Neurotechnology Governance in the EU and US

Elif Kiesow Cortez

2026

TTLF Working Papers

Editors: Siegfried Fina, Mark Lemley, and Roland Vogl

About the TTLF Working Papers

TTLF's Working Paper Series presents original research on technology, and business-related law and policy issues of the European Union and the US. The objective of TTLF's Working Paper Series is to share "work in progress". The authors of the papers are solely responsible for the content of their contributions and may use the citation standards of their home country. The TTLF Working Papers can be found at http://ttlf.stanford.edu. Please also visit this website to learn more about TTLF's mission and activities.

If you should have any questions regarding the TTLF's Working Paper Series, please contact Vienna Law Professor Siegfried Fina, Stanford Law Professor Mark Lemley or Stanford LST Executive Director Roland Vogl at the

Stanford-Vienna Transatlantic Technology Law Forum
http://ttlf.stanford.edu

Stanford Law School
Crown Quadrangle
559 Nathan Abbott Way
Stanford, CA 94305-8610

University of Vienna School of Law
Department of Business Law
Schottenbastei 10-16
1010 Vienna, Austria

About the Author

Dr. Elif Kiesow Cortez has experience in working in policy advisory, project management, and academic roles related to the governance of emerging technologies. Elif collaborates on a research project with the Stanford Institute for Human-Centered Artificial Intelligence. She has been appointed to the advisory boards of projects of leading institutions, including the UNFCCC, IAPP, and the European Research Council. She has also founded and led the AI & LegalTech Lab. Previously, Elif was a John M. Olin Fellow in Law and Economics at Harvard Law School. Her doctoral research at the Institute of Law and Economics at the University of Hamburg in Germany was funded by the German Research Association (DFG). During her doctoral studies, she was a visiting researcher at Harvard Business School and the University of California, Berkeley, School of Law. Elif is an expert on behavioral strategies for effective tech policies that address cooperation problems between public and private actors. She has acquired research grants for tech governance projects commissioned by the Dutch Research Council (NWO) and the US National Science Foundation (NSF). Elif has been a TTLF Affiliate since 2020.

General Note about the Content

The opinions expressed in this paper are those of the author and not necessarily those of the Transatlantic Technology Law Forum or any of its partner institutions, or the sponsors of this research project.

Suggested Citation

This TTLF Working Paper should be cited as:
Elif Kiesow Cortez, Legal and Technical Pathways for Neurotechnology Governance in the EU and US, Stanford-Vienna TTLF Working Paper No. 161, http://ttlf.stanford.edu.

Copyright

© 2026 Elif Kiesow Cortez

Abstract

Advancements in neurotechnology pose a significant challenge to the legal frameworks that protect mental privacy. Until recently, devices capable of recording, interpreting, or modulating neural activity were confined to clinical and laboratory settings. Now, they are being deployed across consumer, educational, and workplace environments. The neural data generated by these systems can reveal information about emotional and cognitive states and complicate the already controversial question of what constitutes personal data. Existing privacy and data protection frameworks may be both conceptually and structurally inadequate when confronted with technology that can address itself directly to cognitive activity. Drawing on a comparison of the United States' and the European Union's regulatory approaches, this paper explores how legal standards can incorporate technical privacy-enhancing mechanisms, such as encryption, differential privacy, and local processing. This integration aims to operationalize the protection of neural information within specific sectoral contexts. By integrating these legal and technical considerations, the paper articulates a framework for neural data governance that balances innovation incentives with the safeguarding of mental privacy. The analysis situates this framework within the broader legal context of recognizing cognitive rights, sectoral adaptation, and institutional oversight.

Table of Contents

I. Introduction 3

1. The first judicial application of neurorights: Girardi v. Emotiv Inc. 6

II. The Exceptionalism Debate and Its Shared Premise 10

1. The case for novelty 10

2. The case against 11

3. Synthesis of the positions 13

III. The Genetic Precedent 14

IV. The European Framework 20

1. The GDPR: general personal data, special categories, and privacy by design 20

2. The Artificial Intelligence Act and inferential uses 24

3. Medical-device law, consumer neurotechnology, and the wellness-medical gray zone .. 27

4. European Health Data Space, research, and institutional governance 30

V. The American Framework 32

1. HIPAA: protection within a limited institutional perimeter 32

2. FTC authority, FDA medical-device oversight, and the Common Rule 34

3. The State-level neural-data laws 36

4. Federal proposals and constitutional dimensions 39

5. Carpenter v. United States and revealing third-party records 40

6. Empirical audit evidence 41

VI. Biometric and Genetic Analogies 42

VII. Neural Decoding Capabilities 45

1. Semantic decoding and the cooperation constraint 45

2. Inner speech and representational safeguards 47

3. Identifiability 48

4. The category problem 49

VIII. Privacy-Enhancing Technologies and Their Legal Significance 50

1. Cryptographic approaches 50

2. Differential privacy and synthetic data 51

3. Local processing and perturbation 52

4. The need for a synthesis approach 53

IX. Translating Legal Standards into Technical Requirements 56

1. The formalization method 56

2. Standards & Specification 57

3. A graded framework 58

X. Sector-Specific Application 62

1. Clinical neurotechnology 62

2. Employment 63

3. Consumer devices 64

4. Research 65

5. Education 66

XI. Towards Transatlantic Harmonization 67

XIII. Conclusion 72

References 75

I. Introduction

Neurotechnology has expanded from a predominantly clinical and laboratory domain into a considerably more heterogeneous ecosystem, one that now encompasses medical implants, research platforms, consumer electroencephalography headsets, neuroadaptive interfaces, attention and fatigue monitors, and artificial intelligence systems capable of classifying or decoding features of nervous-system activity.[1] [2] This expansion brought changes in the associated legal problems and therefore created a need for further discussions on how to address potential gaps regarding the protection of neural information. Within a hospital, an electroencephalographic (EEG) recording is ordinarily embedded within a professional relationship and situated within sector-specific health regulation, so that the applicable legal obligations follow more or less predictably from the institutional setting in which the recording was made. Where the same or a functionally equivalent signal is generated by a consumer device, a workplace monitoring tool, or an educational platform, it may be produced entirely outside those protections, and may subsequently be combined with cloud analytics, advertising identifiers, behavioral telemetry, or machine-learning models in ways that the person from whom the signal originated is unlikely to have anticipated.[3] This requires a regulatory governance approach that goes beyond the device considered in isolation, and expands into the interaction among the signal, the inference drawn from it, the purpose for which it is processed, the sector in which processing occurs, the identity of the actor conducting that processing, and the technical architecture within which it is

[1] 'The Protection of Mental Privacy in the Area of Neuroscience - Societal, Legal and Ethical Challenges | Think Tank | European Parliament' <https://www.europarl.europa.eu/thinktank/en/document/EPRS_STU(2024)757807> accessed 14 September 2026.
[2] 'TechDispatch #1/2024 - Neurodata | European Data Protection Supervisor' <https://www.edps.europa.eu/data-protection/our-work/publications/techdispatch/2024-06-03-techdispatch-12024-neurodata> accessed 14 September 2026.
[3] Sjors Ligthart and others, 'Minding Rights: Mapping Ethical and Legal Foundations of "Neurorights"' (2023) 32 Cambridge Quarterly of Healthcare Ethics 461 <https://doi.org/10.1017/S0963180123000245>.

embedded.[4] [5] [6] International bodies addressing the field have converged upon a comparable framing.[7] [8]

The scientific capability underlying these concerns would need to be characterized in a granular manner in order to avoid either sensationalism or dismissal. It is established that non-invasive functional magnetic resonance imaging (fMRI) studies have reconstructed semantic content under cooperative and subject-specific conditions[9], intracortical systems have achieved high-performance decoding of attempted speech[10] [11], that research employing wearable electroencephalography and magnetoencephalography has improved sentence and stimulus classification[12], and that long-term independent use of implanted brain-computer interfaces within the home environment[13] has moved the technology beyond the confines of short experimental sessions.

[4] Marcello Ienca and Roberto Andorno, 'Towards New Human Rights in the Age of Neuroscience and Neurotechnology' (2017) 13 Life Sciences, Society and Policy 5 <https://doi.org/10.1186/s40504-017-0050-1>.
[5] Anna Wexler and Peter B Reiner, 'Oversight of Direct-to-Consumer Neurotechnologies' (2019) 363 Science 234 <https://doi.org/10.1126/science.aav0223>.
[6] Nita Farahany, 'The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology' [2023] Faculty Books.
[7] OECD, 'Recommendation of the Council on Responsible Innovation in Neurotechnology' (12 November 2019) <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0457> accessed 14 September 2026.
[8] UNESCO, 'Recommendation on the Ethics of Neurotechnology - Legal Affairs' (11 November 2025) <https://www.unesco.org/en/legal-affairs/recommendation-ethics-neurotechnology> accessed 14 September 2026.
[9] Jerry Tang and others, 'Semantic Reconstruction of Continuous Language from Non-Invasive Brain Recordings' (2023) 26 Nature Neuroscience 858 <https://doi.org/10.1038/s41593-023-01304-9>.
[10] Francis R Willett and others, 'A High-Performance Speech Neuroprosthesis' (2023) 620 Nature 1031 <https://doi.org/10.1038/s41586-023-06377-x>.
[11] Erin M Kunz and others, 'Inner Speech in Motor Cortex and Implications for Speech Neuroprostheses' (2025) 188 Cell 4658 <https://doi.org/10.1016/j.cell.2025.06.015>.
[12] Alexandre Défossez and others, 'Decoding Speech Perception from Non-Invasive Brain Recordings' (2023) 5 Nature Machine Intelligence 1097 <https://doi.org/10.1038/s42256-023-00714-5>.
[13] Nicholas S Card and others, 'Long-Term Independent Use of an Intracortical Brain–Computer Interface for Speech and Cursor Control' (2026) 32 Nature Medicine 2504 <https://doi.org/10.1038/s41591-026-04414-6>.

The clinical-translation pipeline[14] underlying these results has itself become substantial, building upon two decades of communication-oriented research[15] in the field. At the same time, it bears emphasis that contemporary systems remain constrained by requirements of calibration, by limitations of signal quality, by the design of the decoding task, by contextual factors, and by the cooperation of the user. The legal significance of these developments therefore lies less in any claim that present-day devices are capable of reading arbitrary thoughts than in what may be termed the temporal instability of neural data: a recording which supports only modest inferences at the moment of its acquisition may become considerably more revealing once improved models, auxiliary datasets, or AI-enhanced superior decoding methods become available. This relevant instability problem is already discernible in findings indicating that non-clinical electroencephalography can track politically consequential attitudes[16], and in the parallel literature considering whether neural measures may be advanced as evidence of memory in legal proceedings[17].

This paper firstly critically reviews the scholarly literature on neurorights and situates that debate within the older controversy over genetic exceptionalism, contending that both the proponents and the critics of a special legal category for neural data have left underappreciated how classification is supposed to translate into protection. It then turns to a comparative examination of the two regulatory traditions with which it is principally concerned, tracing the European approach through

[14] K Michelle Patrick-Krueger, Ian Burkhart and Jose L Contreras-Vidal, 'The State of Clinical Trials of Implantable Brain–Computer Interfaces' (2025) 3 Nature Reviews Bioengineering 50 <https://doi.org/10.1038/s44222-024-00239-5>.
[15] Niels Birbaumer, 'Breaking the Silence: Brain–Computer Interfaces (BCI) for Communication and Motor Control' (2006) 43 Psychophysiology 517 <https://doi.org/10.1111/j.1469-8986.2006.00456.x>.
[16] Giulia Galli and others, 'Early EEG Responses to Pre-Electoral Survey Items Reflect Political Attitudes and Predict Voting Behavior' (2021) 11 Scientific Reports 18692 <https://doi.org/10.1038/s41598-021-96193-y>.
[17] Emily RD Murphy and Jesse Rissman, 'Evidence of Memory from Brain Data' (2020) 7 Journal of Law and the Biosciences lsaa078 <https://doi.org/10.1093/jlb/lsaa078>.

the General Data Protection Regulation, the Artificial Intelligence Act, the Medical Devices Regulation, and the European Health Data Space, and the American approach through the sectoral and actor-specific regimes of the Health Insurance Portability and Accountability Act, medical-device oversight, the Common Rule, and the Federal Trade Commission Act, together with the expanding patchwork of state legislation. Against that legal background the paper assesses what neural decoding has actually been demonstrated to accomplish and which legal implications might come out of it.

From this foundation the paper develops its central claim, which is that the protection promised by any legal category becomes operational only where the abstract standards of security, minimization, and data protection by design are rendered technically reviewable. It accordingly surveys the principal privacy-enhancing technologies, among them encryption, secure computation, differential privacy, local processing, and representation-level safeguards, and considers how each might be attached to a specific legal requirement rather than left to the unstructured discretion of the controller. Drawing these strands together, the paper articulates a graded governance framework keyed to inferential risk rather than to recording modality, applies that framework across five sectors in which the consent mechanism applies differently, and closes by addressing questions of institutional oversight, soft law, and transatlantic harmonization, and addresses potential objections to which the framework is exposed.

1. The first judicial application of neurorights: Girardi v. Emotiv Inc.

The first judicial decision in which a claim framed in terms of neurorights was adjudicated in Chile. In *Girardi v. Emotiv Inc.*, the former Chilean senator Guido Girardi purchased and used the Insight electroencephalography headset manufactured by Emotiv, uploaded recordings of his neural activity to the company's platform, and subsequently objected to the company's retention of

those recordings and to its contemplated use of them for research purposes.[18] He sought constitutional protection upon the ground that the device, together with the associated data practices, endangered his privacy and his physical and psychological integrity.[19] The matter reached the Supreme Court of Chile following the rejection of the claim at first instance. It is important to note that the dispute originated from an ordinary consumer transaction, in which an individual voluntarily acquired a commercial device and accepted the associated privacy policy before subsequently challenging the downstream use and retention of the resulting brain data.[20]

On 9 August 2023 the Supreme Court granted the protection sought. In doing so it emphasized the constitutional protection afforded to physical and psychological integrity and to privacy, the requirement of specific consent where information is to be used for purposes of scientific research, and the degree of precaution appropriate to technologies capable of revealing aspects of the person that earlier legal arrangements had not contemplated.[21] The Court directed the relevant health and customs authorities to examine the device under the applicable regulatory regime, and ordered Emotiv to delete the information it had stored in connection with Girardi's use of the product while declining the proposition that aggregation or anonymization automatically converts such recordings into freely usable statistical material where the consent legally required in respect of the new purpose is absent.[22] The judgment is significant because it demonstrates that the first neurorights case to reach a supreme court leaned on an in depth analysis of familiar questions

[18] 'Supreme Court of Chile, Role 105.065-2023, Judgment of 9-8-2023' <https://www.doe.cl/alerta/11082023/20230811001>.
[19] María Isabel Cornejo-Plaza, Roberto Cippitani and Vincenzo Pasquino, 'Chilean Supreme Court Ruling on the Protection of Brain Activity: Neurorights, Personal Data Protection, and Neurodata' (2024) 15 Frontiers in Psychology 1330439 <https://doi.org/10.3389/fpsyg.2024.1330439>.
[20] Equipo Actualidad Jurídica, 'Corte Suprema ordenó eliminar información recogida por dispositivo de monitoreo de actividad cerebral del recurrente' (*DOE | Actualidad Jurídica*, 11 August 2023) <https://actualidadjuridica.doe.cl/corte-suprema-ordeno-eliminar-informacion-recogida-por-dispositivo-de-monitoreo-de-actividad-cerebral-del-recurrente/> accessed 14 September 2026.
[21] Cornejo-Plaza, Cippitani and Pasquino (n 19).
[22] 'Supreme Court of Chile, Role 105.065-2023, Judgment of 9-8-2023' (n 18).

concerning consent, purpose limitation, regulatory authorization, retention, and deletion. The constitutional amendment for protecting neurorights adopted in Chile supplies the background against which the decision was rendered. [23] [24]

The scholarly literature concerning mental privacy encompasses proposals[25] [26] for the recognition of cognitive liberty, mental integrity, psychological continuity, and a distinct right to mental privacy. There has been a critical literature[27] [28] directed at what its proponents characterize as neuroexceptionalism and rights inflation, including scholarly debate[29] [30] on existing right to freedom of thought, properly interpreted, already extends to brain data, questioning the necessity for the recognition of novel rights. Related scholarship has extended the analysis to mental integrity and to the regulatory landscape considered more broadly[31] [32], and scholars have sought to establish the minimum shared understanding on the issue.[33] It is nevertheless the contention of this paper that both the pro-rights and the anti-exceptionalist positions have tended to leave the central engineering question underspecified. Once the law has determined that a given neural data

[23] 'Chile: Pioneering the Protection of Neurorights | The UNESCO Courier' <https://courier.unesco.org/en/articles/chile-pioneering-protection-neurorights> accessed 14 September 2026.
[24] Sergio Ruiz and others, 'Neurorights in the Constitution: From Neurotechnology to Ethics and Politics' (2024) 379 Philosophical Transactions of the Royal Society B: Biological Sciences 20230098 <https://doi.org/10.1098/rstb.2023.0098>.
[25] Ienca and Andorno (n 4).
[26] Ligthart and others (n 3).
[27] Jan Christoph Bublitz, 'Novel Neurorights: From Nonsense to Substance' (2022) 15 Neuroethics 7 <https://doi.org/10.1007/s12152-022-09481-3>.
[28] Daniel Susser and Laura Y Cabrera, 'Brain Data in Context: Are New Rights the Way to Mental and Brain Privacy?' (2024) 15 AJOB Neuroscience 122 <https://doi.org/10.1080/21507740.2023.2188275>.
[29] Nora Hertz, 'Neurorights – Do We Need New Human Rights? A Reconsideration of the Right to Freedom of Thought' (2022) 16 Neuroethics 5 <https://doi.org/10.1007/s12152-022-09511-0>.
[30] 'Protecting the Mental Realm: What Does Human Rights Law Bring to the Table? - Timo Istace, 2023' <https://journals.sagepub.com/doi/10.1177/09240519231211823> accessed 14 September 2026.
[31] Vera Tesink and others, 'Right to Mental Integrity and Neurotechnologies: Implications of the Extended Mind Thesis' (2024) 50 Journal of Medical Ethics 656 <https://doi.org/10.1136/jme-2023-109645>.
[32] Łukasz Szoszkiewicz and Rafael Yuste, 'Mental Privacy: Navigating Risks, Rights and Regulation : Advances in Neuroscience Challenge Contemporary Legal Frameworks to Protect Mental Privacy' (2025) 26 EMBO reports 3469 <https://doi.org/10.1038/s44319-025-00505-6>.
[33] Ligthart and others (n 3).

flow is protected, it would be necessary to determine what precisely must a controller do in order to render that protection effective. The analysis accordingly treats the recognition of cognitive rights as one layer of governance among several, while directing its principal attention to the translation of existing duties, lawfulness, minimization, confidentiality, purpose limitation, security, impact assessment, and accountability, into technical requirements capable of being audited.

The paper proceeds with providing a transatlantic overview of relevant regulatory efforts in order to shape its exploratory neurotechnology governance framework. The European Union starts from horizontal data-protection and fundamental-rights contexts, to which it has added technology-focused, and sectoral instruments including the Artificial Intelligence Act, the Medical Devices Regulation, and the European Health Data Space. The United States, by contrast, begins from sectoral and actor-specific regimes such as the Health Insurance Portability and Accountability Act, medical-device regulation administered by the Food and Drug Administration, the Common Rule, the Federal Trade Commission Act, and an expanding body of state legislation[34] with the consequence that functionally identical neural information receives materially different protection depending upon the identity of the party which collected it and the purpose for which it was collected.[35]

[34] This includes also five state statutes which constituted the first wave of American neural-data legislation, namely those of Colorado, California, Montana, Connecticut and Vermont which enacted a defining neural data in June 2026, the principal provisions of which take effect in 2028.
[35] Elif Kiesow Cortez, 'Neurotechnology and Privacy: A Comparative Analysis of US and EU Approaches' (TTLF Working Papers, Stanford-Vienna Transatlantic Technology Law Forum 2025).

II. The Exceptionalism Debate and Its Shared Premise

This Part examines the principal scholarly dispute concerning the appropriate legal characterisation of neural data. It proceeds in three sections. The first sets out the case for recognising neural data as a novel category warranting novel protection; the second sets out the principal objections to that position; and the third identifies a premise which both camps share and which neither has adequately examined.

1. The case for novelty

The modern literature on neurorights takes as its point of departure the proposal for the recognition of four new human rights, namely cognitive liberty, mental privacy, mental integrity, and psychological continuity.[36] Scholars proceed from the observation that the existing human rights instruments were drafted upon the assumption that the interior of the mind is, as a practical matter, inaccessible to external observation, and that this assumption is progressively being eroded by the development of devices capable of recording and, increasingly, of modulating neural activity.[37] Mental privacy extends to any brain information recorded by a neurodevice and does so irrespective of whether that information has yet been decoded into intelligible content[38] and thus this also addresses the temporal instability identified in Part I of this paper.

The proposal was extended by a multidisciplinary team of scholars, who identified privacy and consent, agency and identity, augmentation, and bias as the four priority areas for the governance of neurotechnology.[39] This broader intervention did not confine itself to the vocabulary of rights,

[36] Ienca and Andorno (n 4).
[37] ibid.
[38] Ligthart and others (n 3).
[39] Rafael Yuste and others, 'Four Ethical Priorities for Neurotechnologies and AI' (2017) 551 Nature 159 <https://doi.org/10.1038/551159a>.

instead it proposed a series of technical and default-setting responses, including opt-out defaults in respect of the sharing of neural data and the deployment of differential privacy, federated computation, and distributed-ledger techniques.[40]

The importance of a systematic treatment of neurorights was further elaborated[41], and other contributions in the recent years argued that all brain-derived data should be by default classified as sensitive health data, with medical-grade regulation applied irrespective of whether the collecting device is clinical or consumer in character.[42] This position within the debate has proved influential in practice: it underpins the Chilean constitutional amendment[43], it informed the drafting of the California neural data provision[44], and it constitutes the foundation of the advocacy programme pursued by the NeuroRights Foundation.[45]

2. The case against

The broader approach to define neurorights also encountered criticism[46], on three accounts: (1) the proposals might contribute to 'rights inflationism', that is, the proliferation of claimed human rights to a point at which the currency of such claims is progressively depreciated, (2) rest upon neuroexceptionalism and neuroessentialism, being the largely unargued assumptions that brain data is categorically distinct from other forms of data and that the brain constitutes the privileged locus of personhood, (3) none of the individual proposals satisfies the quality criteria established

[40] ibid.
[41] Marcello Ienca, 'On Neurorights' (2021) 15 Frontiers in Human Neuroscience <https://doi.org/10.3389/fnhum.2021.701258> accessed 15 September 2026.
[42] Rafael Yuste, 'Advocating for Neurodata Privacy and Neurotechnology Regulation' (2023) 18 Nature Protocols 2869 <https://doi.org/10.1038/s41596-023-00873-0>.
[43] Sergio Ruiz and others, 'Neurorights in the Constitution: From Neurotechnology to Ethics and Politics' (2024) 379 Philosophical Transactions of the Royal Society B: Biological Sciences 20230098 <https://doi.org/10.1098/rstb.2023.0098>.
[44] (California S.B. 1223)
[45] Genser, Jared, Damianos, Stephen and Yuste, Rafael, 'Safeguarding Brain Data: Assessing the Privacy Practices of Consumer Neurotechnology Companies' (Neurorights Foundation 2024).
[46] Bublitz (n 27).

for the recognition of a new human right, which require among other matters that the interest in question be insufficiently protected by the rights already recognized.

A distinct though converging argument[47] also compares the proposed right to mental self-determination with the existing right to freedom of thought under Article 9 of the European Convention on Human Rights and Article 18 of the International Covenant on Civil and Political Rights, debating that freedom of thought may coherently be read as protecting mental processes and brain data in their entirety, with the consequence that an evolving interpretation of the existing right is more convincing than the recognition of a new one. A parallel case for interpretive evolution rather than novel rights has been also debated further in the literature[48] [49] and constitutes of claims concerning the sufficiency of existing law rather than a sole claim concerning the sensitivity of neural data, and it therefore exhibits the same structure as the argument which this paper advances in relation to Articles 25 and 32 of the General Data Protection Regulation.

Some scholars mount a different though complementary challenge from within American privacy theory.[50] Applying Nissenbaum's framework of contextual integrity, they argue that the privacy risks presently posed by neurotechnology are not, at least for the time being, different in kind from those raised by genomic sequencing, location tracking, or behavioural advertising.[51] Their prescription is not that protection should be abandoned but rather that it should be relocated: instead of codifying a freestanding right to mental privacy, regulators should ask what informational norms govern particular contexts, healthcare, criminal justice, consumer marketing,

[47] Hertz (n 29).
[48] Sjors Ligthart, 'Freedom of Thought in Europe: Do Advances in "Brain-Reading" Technology Call for Revision?' (2020) 7 Journal of Law and the Biosciences lsaa048 <https://doi.org/10.1093/jlb/lsaa048>.
[49] Timo Istace, 'Establishing Neurorights: New Rights versus Derived Rights' (2025) 17 Journal of Human Rights Practice 121 <https://doi.org/10.1093/jhuman/huae042>.
[50] Daniel Susser and Laura Y Cabrera, 'Brain Data in Context: Are New Rights the Way to Mental and Brain Privacy?' (2024) 15 AJOB Neuroscience 122 <https://doi.org/10.1080/21507740.2023.2188275>.
[51] ibid.

and to enforce those norms accordingly.[52] This constitutes a directly sector-specific recommendation, and it accords closely with the approach developed in Part X below.

Multidisciplinary work also seeks to establish a minimalist shared understanding of mental privacy, mental integrity, and cognitive liberty sufficient to permit scholarly and legislative debate to proceed without the participants talking past one another.[53]

3. Synthesis of the positions

The debate has been conducted largely as a question concerning categories: does neural data belong within a special legal box? Both camps, however, share an assumption which merits closer scrutiny than it has hitherto received, namely the assumption that the protective work is performed at the point of classification. The exceptionalists argue that placing neural data within a special category will produce heightened protection[54] [55], while the critics argue that it will not produce protection commensurate with the disruption occasioned and that the existing categories are sufficient.[56] [57] Neither position interrogates the mechanism by which classification is supposed to translate into protection, and it is that mechanism which forms the subject of the present paper.

Article 32 of the General Data Protection Regulation requires security appropriate to the risk, offering pseudonymization and encryption by way of example rather than by way of requirement.[58] Article 25 requires data protection by design and by default, taking into account the state of the

[52] ibid.
[53] Ligthart and others (n 3).
[54] Yuste (n 42).
[55] Marcello Ienca and others, 'Towards a Governance Framework for Brain Data' (2022) 15 Neuroethics 20 <https://doi.org/10.1007/s12152-022-09498-8>.
[56] Bublitz (n 27).
[57] Susser and Cabrera (n 28).
[58] (Regulation (EU) 2016/679, art. 32)

art.[59] [60] The American state statutes generally incorporate a reasonable-security obligation drawn from the underlying consumer privacy legislation.[61] In each instance the legal category determines that a heightened standard is applicable, but the content of that heightened standard is left to be supplied by the regulated entity, subject only to enforcement after the fact.[62]

In respect of the majority of data types this delegation is defensible. The technical measures appropriate to a customer database are broadly understood within the relevant professional community, and the distance between best and worst practice is not so great as to render the standard meaningless. In the case of neural data the delegation is considerably less defensible, and for three reasons. First, the relevant technical measures are neither standardized nor widely known.[63] Second, the distance between best and worst practice is very substantial. Third, and most importantly, the available empirical evidence concerning present governance practice, which is examined in Part V below, indicates that the modal practice lies close to the worst case.[64] The category has been established; its content remains to be supplied. It is from this observation that the argument of this paper proceeds.

III. The Genetic Precedent

If the question to be answered is whether the creation of a special legal category in respect of a biologically derived data type produces meaningful protection in practice, there exists one

[59] ibid.
[60] European Data Protection Board, 'Guidelines 4/2019 on Article 25 Data Protection by Design and by Default' (2020).
[61] 'The "Neural Data" Goldilocks Problem: Defining "Neural Data" in U.S. State Privacy Laws' (*https://fpf.org/*) <https://fpf.org/blog/the-neural-data-goldilocks-problem-defining-neural-data-in-u-s-state-privacy-laws/> accessed 15 September 2026.
[62] Ari Ezra Waldman, 'Data Protection by Design? A Critique of Article 25 of the GDPR' [2021] Cornell International Law Journal, Vol. 53, 2021.
[63] Kun Xia and others, 'Privacy-Preserving Brain–Computer Interfaces: A Systematic Review' (2023) 10 IEEE Transactions on Computational Social Systems 2312 <https://doi.org/10.1109/TCSS.2022.3184818>.
[64] Genser, Jared, Damianos, Stephen and Yuste, Rafael (n 45).

substantial body of comparative experience upon which to draw. If the question to be answered is whether the creation of a special legal category in respect of a biologically derived data type produces meaningful protection in practice, there exists one substantial body of comparative experience upon which to draw. The earlier controversy over genetic exceptionalism has been invoked by several scholars within the neurorights debate, whether as a cautionary parallel or as evidence that a special legal category tends to disappoint, and the suggestion that the concerns raised by neural data are continuous with those long familiar from genomics has been advanced for many years. This analysis focuses on the specific trajectory by which the genetics field moved from exceptionalism to genomic contextualism, identifies the definitional-contest failure mode that survived that transition, and carries that failure mode through to the operative text of the recent neural data statutes and, in the parts which follow, to the technical question of what renders a legal standard on neurotechnology reviewable.

Between 1990 and 2008 nearly every American state enacted legislation prohibiting genetic discrimination in health insurance, approximately two-thirds did so in respect of employment[65], and Congress ultimately enacted the Genetic Information Nondiscrimination Act.[66] The European Union, for its part, placed genetic data within the special categories established by the General Data Protection Regulation.[67] The debate concerning neural data is, in its structure, a similar debate concerning genetic exceptionalism, and the resemblance extends even to the terminology employed. Murray introduced the expression "genetic exceptionalism" in 1997, to denote the claim

[65] Mark A Rothstein, 'Genetic Exceptionalism and Legislative Pragmatism' (2007) 35 Journal of Law, Medicine & Ethics 59 <https://doi.org/10.1111/j.1748-720X.2007.00154.x>.
[66] Yann Joly and others, 'Looking Beyond GINA: Policy Approaches to Address Genetic Discrimination' (2020) 21 Annual Review of Genomics and Human Genetics 491 <https://doi.org/10.1146/annurev-genom-111119-011436>.
[67] (Regulation (EU) 2016/679, art. 9).

that genetic information is so distinct in concept, in practical implication, and in moral import as to warrant being singled out from other health-related information[68] [69] and the charge of neuroexceptionalism advanced by Bublitz is its lineal descendant.[70]

As to efficacy, scholars examined the landscape of genetic privacy across federal statutes and regulations, including both the Health Insurance Portability and Accountability Act and the Genetic Information Nondiscrimination Act, and concluded that few if any applicable legal doctrines or enactments furnish adequate protection or meaningful control to the individuals whose information is at stake.[71] That conclusion is the more striking for having been reached by Rothstein a second time, more than a decade after his earlier and considerably more preliminary assessment[72], and upon a far larger evidentiary base, since the intervening period saw the arrival of direct-to-consumer sequencing, of consumer genealogy databases, and of forensic genetic genealogy, none of which the original statutes contemplated. The mechanism of the shortfall is worth stating precisely, because it transfers. Genetic-specific statutes required a determination that particular information was genetic in origin, a requirement which invited definitional litigation and which permitted functionally equivalent information derived from family history, from medical records, or from biomarkers to fall outside the protected category.[73] They addressed discrimination rather than collection, with the consequence that the underlying data flows remained intact.[74] And

[68] Thomas H Murray, 'Is Genetic Exceptionalism Past Its Sell-By Date? On Genomic Diaries, Context, and Content' (2019) 19 The American Journal of Bioethics 13 <https://doi.org/10.1080/15265161.2018.1552038>.
[69] Thomas H Murray, 'Genetic Exceptionalism and 'Future Diaries' : Is Genetic Information Different from Other Medical Information?' [1997] Genetic Secrets : Protecting Privacy and Confidentiality in the Genetic Era 60. cited by Murray (n 68).
[70] Bublitz (n 27).
[71] Ellen Wright Clayton and others, 'The Law of Genetic Privacy: Applications, Implications, and Limitations' (2019) 6 Journal of Law and the Biosciences 1 <https://doi.org/10.1093/jlb/lsz007>.
[72] Rothstein (n 65).
[73] Clayton and others (n 71).
[74] ibid.

their coverage was incomplete in ways which the category itself did not disclose: the Genetic Information Nondiscrimination Act reaches health insurance and employment but not life, disability, or long-term care insurance, which are the contexts in which the concern regarding genetic discrimination has proved most persistent.[75] [76]

Scholars reporting from the International Genetic Discrimination Observatory established for the purpose, found that genetic discrimination continues to cast a substantial shadow notwithstanding two decades of category-based legislation[77], and the establishment of a standing observatory is itself an admission of a kind.[78] It is the response of a field which has concluded that the problem requires continuous empirical monitoring rather than a further refinement of statutory definitions. As to the conceptual apparatus, the development is more consequential still for present purposes, because the genetics community has not merely criticized exceptionalism but has largely abandoned it and adopted a successor framework. Scholars also proposed genomic contextualism in its place, arguing that genomic information is in certain respects continuous with other health information and in certain respects distinct, and that the useful question is accordingly not whether it belongs to a special class but what treatment is appropriate within a given context of use.[79] Murray, responding to that proposal, agreed that the term he had introduced should be retired, observing that it had been pulled in so many and such contradictory directions that its meaning

[75] Joly and others (n 66).
[76] Robert C Green, Denise Lautenbach and Amy L McGuire, 'GINA, Genetic Discrimination, and Genomic Medicine' (2015) 372 The New England Journal of Medicine 397 <https://doi.org/10.1056/NEJMp1404776>.
[77] Yann Joly and Gratien Dalpe, 'Genetic Discrimination Still Casts a Large Shadow in 2022' (2022) 30 European Journal of Human Genetics 1320 <https://doi.org/10.1038/s41431-022-01194-8>.
[78] Yann Joly and others, 'Establishing the International Genetic Discrimination Observatory' (2020) 52 Nature Genetics 466 <https://doi.org/10.1038/s41588-020-0606-5>.
[79] Nanibaa' A Garrison and others, 'Genomic Contextualism: Shifting the Rhetoric of Genetic Exceptionalism' (2019) 19 The American Journal of Bioethics 51 <https://doi.org/10.1080/15265161.2018.1544304>.

and usefulness had dissipated, and adding the caution that attention to context should apply generally rather than being reserved for whichever data type is currently under discussion.[80]

Three consequences follow for the argument of this paper, and they are more pointed than the cautionary reading which the earlier literature alone supports. The first is that the neural data debate is not merely rehearsing an argument which happened to fail. It is rehearsing an argument which the relevant literature has already concluded, and concluded against the exceptionalist position, in favor of a contextual one. That is a heavier burden for the proponents of a special neural category to discharge than the mere observation that genetic legislation underperformed, and it is a burden which the current literature does not appear to have registered. It is also, it may be noted, the same conclusion at which scholars arrive by way of contextual integrity, and their argument is strengthened rather than weakened by the observation that the adjacent field reached it independently and some years earlier.[81]

The second is that the residual difficulty, once the definitional question is set aside, turns out to lie in scope of coverage and in enforcement architecture rather than in the boundary of the category. The Genetic Information Nondiscrimination Act does not fail because "genetic information" was badly defined; it falls short because it does not reach the insurance markets in which the apprehended harm is concentrated, and because the remedies it provides are difficult to invoke.[82][83] That is precisely the distinction which this paper draws between a category, which is a

[80] Murray (n 68).
[81] Susser and Cabrera (n 28).
[82] Joly and others (n 66).
[83] Clayton and others (n 71).

jurisdictional device determining which rules apply, and the rules themselves, whose content determines whether anything follows from the classification.

The third is that each of the failure modes so identified is discernible in embryonic form within the existing neural data statutes, and can now be named with some precision. The Colorado provision[84] regulates biological data only where processed for purposes of identification, a qualifier which invites the definitional contest that Rothstein described[85] and which Part VII will address regarding the available evidence concerning neural identifiability.

The Connecticut definition[86] extends only to activity of the central nervous system, thereby excluding peripheral measurement. The California definition[87] excludes information inferred from nonneural sources, which is the location in which a considerable proportion of the inferential risk in fact resides.[88] The Montana approach[89], which proceeds by amending a genetic privacy statute, is straightforward concerning the analogy it draws.

The lesson to be drawn is accordingly not that special categories are without utility, nor even the weaker proposition that they have historically disappointed. It is that the field which first constructed such a category has since dismantled it in favour of contextual assessment, and that the difficulties which survived the category were difficulties of coverage and of enforcement. Neural exceptionalism risks failing in the same manner, and for a reason which is now considerably

[84] (Colorado H.B. 24-1058).
[85] Rothstein (n 65).
[86] (Connecticut Public Act 25-113).
[87] (California S.B. 1223).
[88] Patrick Magee, Marcello Ienca and Nita Farahany, 'Beyond Neural Data: Cognitive Biometrics and Mental Privacy' (2024) 112 Neuron 3017 <https://doi.org/10.1016/j.neuron.2024.09.004>.
[89] (Montana S.B. 163).

better documented than when it was first identified[90]: it triggers security and design obligations the content of which no one has yet specified.[91] [92] The response, in both instances, is to attend to the operative rules rather than to the boundary of the category.

IV. The European Framework

1. The GDPR: general personal data, special categories, and privacy by design

The General Data Protection Regulation, Regulation (EU) 2016/679, constitutes the horizontal framework governing the processing of personal data within the European Union. It applies to information relating to an identified or identifiable natural person and, by virtue of Article 3, may also extend to certain controllers and processors established outside the Union where they offer goods or services to persons within the Union or monitor their behaviour therein. It is significant that the Regulation contains no standalone statutory category denominated "neural data". The absence is consequential, because neural information must first be mapped onto the Regulation's existing concepts of personal data, data concerning health, biometric data, or other special-category information before the enhanced regime established by Article 9 becomes applicable.[93] [94] Neural data generated in the course of diagnosis or treatment will ordinarily constitute health data; neural patterns processed for the purpose of unique identification may qualify as biometric

[90] Sonia M Suter, *The Allure and Peril of Genetics Exceptionalism: Do We Need Special Genetics Legislation?* (2001).
[91] Waldman (n 62).
[92] Ira S Rubinstein and Nathaniel Good, 'The Trouble with Article 25 (and How to Fix It): The Future of Data Protection by Design and Default' (2020) 10 International Data Privacy Law 37 <https://doi.org/10.1093/idpl/ipz019>.
[93] Marcello Ienca and Gianclaudio Malgieri, 'Mental Data Protection and the GDPR' (2022) 9 Journal of Law and the Biosciences lsac006 <https://doi.org/10.1093/jlb/lsac006>.
[94] 'TechDispatch #1/2024 - Neurodata | European Data Protection Supervisor' (n 2).

data; and consumer signals employed for entertainment or device control may remain ordinary personal data notwithstanding that they are highly revealing.[95] [96]

Several GDPR provisions are of particular importance (arts 6, 9, 12–22, 25, 32, 35, 36) within the neurotechnological context because they permit the regulation of the entire processing lifecycle: starting with the initial acquisition of a brain signal feature extraction and encompassing collection, model training, inference, sharing, retention, deletion, and model reuse.[97]

Articles 25 and 32 provide the principal bridge between the legal framework and privacy-enhancing technologies, and they accordingly warrant closer examination. Article 25 expressly requires consideration of the state of the art, the cost of implementation, the context of processing, and the associated risk.[98] Article 32 identifies measures such as pseudonymization and encryption and requires ongoing testing and evaluation of security.[99] The provisions are intentionally technology-neutral, and that neutrality is in principle desirable, but it gives rise to a specification problem of some seriousness.[100] [101] A controller may assert that it has adopted appropriate safeguards without thereby being required to explain whether raw electroencephalographic signal leaves the device, whether cloud operators receive plaintext, whether model updates leak participant-level information, whether a differential-privacy budget has been established, or whether an architecture described as local nevertheless emits reconstructable telemetry. It is

[95] Stephen Rainey and others, 'Is the European Data Protection Regulation Sufficient to Deal with Emerging Data Concerns Relating to Neurotechnology?' (2020) 7 Journal of Law and the Biosciences lsaa051 <https://doi.org/10.1093/jlb/lsaa051>.
[96] Ienca and others (n 55).
[97] Rainey and others (n 95).
[98] European Data Protection Board, 'Guidelines 4/2019 on Article 25 Data Protection by Design and by Default' (n 60).
[99] European Data Protection Board, 'Guidelines 01/2025 on Pseudonymisation' (2025).
[100] Waldman (n 62).
[101] Ira S Rubinstein and Nathaniel Good, 'The Trouble with Article 25 (and How to Fix It): The Future of Data Protection by Design and Default' (2020) 10 International Data Privacy Law 37 <https://doi.org/10.1093/idpl/ipz019>.

submitted that the legal standard requires a technical profile capable of rendering such assertions verifiable.[102] [103]

The European Union has to date enacted no neurotechnology-specific legislation. Protection consequently depends upon two questions: whether neural data falls within the existing instruments, and, if it does, what those instruments require of a controller. Upon the first question there exists broad if not universal scholarly agreement. Upon the second there exists very little analysis of any kind, and it is to this deficiency that the present paper is principally addressed.

Rainey et al. (2020) provided the first systematic examination of the matter, tracing brain data through the full processing pipeline from raw signal to operationalized output and testing the concepts employed by the Regulation against each successive stage. Their findings identify three structural difficulties which merit separate statement. First, neural data is frequently generated involuntarily, a circumstance which places considerable strain upon the consent model, since an individual cannot meaningfully consent to the generation of signals over which he or she exercises no control.[104] Second, neural data is highly reprocessable, in the sense that a recording gathered for one purpose retains latent information extractable by analytic methods developed subsequently, with the consequence that purpose limitation must govern not merely the initial use but an indefinite series of future ones.[105] Third, the boundary between medical device and wellness product permits consumer neurotechnology to escape the stricter regulatory classification which

[102] Aloni Cohen and Kobbi Nissim, 'Towards Formalizing the GDPR's Notion of Singling out' (2020) 117 Proceedings of the National Academy of Sciences of the United States of America 8344.
[103] 'NIST Privacy Framework Version 1.0' [2020] NIST <https://www.nist.gov/privacy-framework> accessed 16 September 2026.
[104] Rainey and others (n 95).
[105] ibid.

its clinical equivalents attract.[106] [107] They conclude that all brain recordings might arguably to be treated as sensitive irrespective of the purpose of collection.[108]

Ienca and Malgieri (2022) broaden the frame of analysis considerably by introducing the concept of mental data, which they define as any data processed for the purpose of inferring cognitive, emotional, or conative states. This category is deliberately wider than that of brain data, since it captures affective computing, behavioural inference, and interaction telemetry alongside electroencephalography.[109] Their examination of the lawful bases available under the Regulation leads them to a procedural proposal, namely a dedicated mental data protection impact assessment. That proposal is of considerable significance for present purposes, because an impact assessment is precisely the legal instrument through which technical measures may be rendered concrete and susceptible to review.[110]

Ienca et al. (2022) argues that brain data ought to be treated as a special category under Article 9 alongside genetic data and who then, proceed to specify the consequences which would follow. Special-category status would engage the requirement of data protection by design and by default under Article 25; the requirement of security appropriate to the risk under Article 32, with pseudonymization and encryption named as illustrative measures; the breach notification obligations arising under Articles 33 and 34; and Article 89(2), which permits derogations in respect of scientific research and the application of which to commercially funded research stands

[106] ibid.
[107] Anna Wexler and Peter B Reiner, 'Oversight of Direct-to-Consumer Neurotechnologies' (2019) 363 Science 234 <https://doi.org/10.1126/science.aav0223>.
[108] Rainey and others (n 95).
[109] Ienca and Malgieri (n 93).
[110] ibid.

in need of clarification.[111] In the European context, this mapping constitutes a possible suggestions regarding which precise provisions the technical specification might attach to.

2. The Artificial Intelligence Act and inferential uses

Regulation (EU) 2024/1689, the Artificial Intelligence Act, adds a further layer which is oriented towards uses and systems rather than towards data. Unlike the General Data Protection Regulation, which follows the processing of personal data, the Artificial Intelligence Act classifies practices and systems according to the risk which they are deemed to present. Article 5 prohibits specified manipulative and exploitative practices, and prohibits in addition artificial intelligence systems deployed for the purpose of inferring emotions in workplaces and educational institutions, subject to exceptions relating to medical and safety purposes. The distinction is of consequence within the neurotechnological context: an electroencephalographic signal employed to control a cursor is not regulated in the same manner as an artificial intelligence system which employs electroencephalography to infer the emotional state of an employee, notwithstanding that the underlying sensor may be identical in both cases[112] and the Act therefore demonstrates the value of regulating downstream inference and context.

The Artificial Intelligence Act[113] entered into force in 2024 and the rules governing prohibited practices began to apply in February 2025. By September 2026 its implementation timetable had been modified by Regulation (EU) 2026/1744, the Digital Omnibus on AI[114], which delayed the application of parts of the high-risk regime. Those alterations of timing do not, however, diminish

[111] Ienca and others (n 55).
[112] Ekaterina Muhl, 'The Challenge of Wearable Neurodevices for Workplace Monitoring: An EU Legal Perspective' (2024) 6 Frontiers in Human Dynamics <https://doi.org/10.3389/fhumd.2024.1473893> accessed 16 September 2026.
[113] (Regulation (EU) 2024/1689)
[114] (Regulation (EU) 2026/1744)

the relevance either of the prohibitions already applicable or of the extraterritorial scope of the Act. For providers operating across the Atlantic the practical consequence is that legality under data protection law and legality under artificial intelligence law are cumulative rather than alternative: a neural processing operation may comply in full with the security requirements of the General Data Protection Regulation and nevertheless remain impermissible because the inference itself is prohibited within the particular context. It may be noted that at the April 2026 convening of the European Brain Council considered in Part XI below, the view was raised that the Artificial Intelligence Act may in fact prove more consequential than the General Data Protection Regulation for the governance of neurotechnology.[115]

The European Group on Ethics in Science and New Technologies has identified a structural limitation in the Act which bears directly upon the argument advanced here. The risk classification of the Act is organized around identifiable systems placed on the market for specified intended uses, whereas the distinctive hazard of neuro-AI is that its capabilities accrue cumulatively, through the aggregation of neural data into reusable foundation models and their subsequent recombination across contexts.[116] A device which appears low-risk when considered in isolation, a consumer headset recording only crude signals, with no genuine capacity to infer or to identify, may nevertheless deliver valuable training data to a Brain Foundation Model, and that model may in turn underpin downstream applications which fall within the prohibited or high-risk categories of the Act, including manipulation, the exploitation of vulnerabilities, emotion recognition, or biometric categorization.[117] Because the Act assesses the device at the point of market placement

[115] European Brain Council (EBC), 'Event Highlights: A European Strategy for Neurotechnology' (20 April 2026) <https://www.braincouncil.eu/event-highlights-a-european-strategy-for-neurotechnology/> accessed 16 September 2026.
[116] European Group on Ethics in Science and New Technologies, 'Governing Neuro-AI: Towards an Infrastructure Approach' (2026).
[117] ibid.

rather than the infrastructure across its lifecycle, these cumulative capabilities may escape classification at the moment when intervention is least costly. The Group accordingly argues, and this paper concurs, that application-specific regulation ought to be complemented by oversight at the reusable model and infrastructure layer, requirements as to the provenance and lawful basis of training data, model documentation, and monitoring throughout the lifecycle, rather than being confined to the assessment of individual devices and their declared purposes.[118]

The significance of the prohibition for neurotechnology is considerable, though it is frequently misstated within the secondary literature. The prohibition is framed around emotion inference rather than around neural measurement, with the consequence that an electroencephalography-based attention or affect classifier deployed within a workplace falls within its scope, whereas a neural device employed for purposes of device control does not.[119] [120] This constitutes a use-based rather than a source-based prohibition, and it is submitted that it furnishes a superior model for the regulation of neural data than do the source-based definitions adopted in the American states, precisely because its application does not depend upon resolving the question whether a given signal is neural in origin.[121]

The prohibition raises in addition a question to which Part VII returns, namely whether the inference which it prohibits is scientifically reliable at all. Scholars found that the relationship between facial movements and emotional states is weak, highly variable, and heavily context-dependent, and that these characteristics are sufficient to undermine the validity of automated

[118] ibid.
[119] (Regulation (EU) 2024/1689, art. 5(1)(f))
[120] Muhl (n 112).
[121] Magee, Ienca and Farahany (n 88).

emotion recognition conducted from facial data.[122] Whether the same conclusion extends to neural emotion inference remains unresolved in the present state of the evidence. It is nevertheless worth observing that the possibility that a prohibited capability does not reliably exist alters the character of the harm in question, transforming it from the disclosure of true inner states into the false attribution of inferred ones, which can still cause harm.

This scepticism as to reliability is shared by the European supervisory authorities. In their joint assessment of the Artificial Intelligence Act the European Data Protection Board and the European Data Protection Supervisor considered that the use of artificial intelligence to infer the emotions of a natural person is highly undesirable and ought to be prohibited save in narrowly specified cases such as health or research, and recommended in addition the explicit prohibition of artificial intelligence systems intended for use by law-enforcement authorities as polygraphs or similar tools, the scientific validity of which is not established.[123] The point transfers directly to the neurotechnological setting: where a prohibited inference is also an unreliable one, the regulatory objection rests not upon the accuracy of the intrusion but upon the consequences of acting upon a false attribution, a matter which the European Data Protection Supervisor treats as itself sufficient to render certain uses of neurodata likely unlawful under Union law.[124]

3. Medical-device law, consumer neurotechnology, and the wellness-medical gray zone

The Medical Devices Regulation, Regulation (EU) 2017/745, supplies a third legal layer. Its principal function concerns product safety and performance rather than informational privacy, but

[122] Lisa Feldman Barrett and others, 'Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements' (2019) 20 Psychological Science in the Public Interest 1 <https://doi.org/10.1177/1529100619832930>.
[123] 'TechDispatch #1/2024 - Neurodata | European Data Protection Supervisor' (n 2).
[124] ibid.

it is of particular relevance to neurotechnology because Annex XVI brings certain products having no intended medical purpose within the medical-device framework.[125] Annex XVI includes equipment intended for brain stimulation which applies electrical currents or magnetic or electromagnetic fields penetrating the cranium so as to modify neuronal activity (ibid.). The consequence is that certain non-therapeutic stimulation technologies are subject to medical-device oversight notwithstanding that they are marketed for purposes of enhancement or wellness. Read-only consumer electroencephalography products, by contrast, may remain outside the Regulation where they lack a qualifying medical purpose, with the result that the General Data Protection Regulation, the Artificial Intelligence Act, and general product and consumer law are required to perform a correspondingly greater share of the regulatory work.[126]

A further and independently significant route to heightened protection is available within the existing law, and it does not depend upon the enactment of any neuro-specific category. Following the position taken by the French supervisory authority, a distinction may be drawn between data which are health data by nature and data which are health data by destination, the latter being ordinary data which acquire the character of health data by reason of the purpose for which, or the inferences to which, they may be put; a wellbeing application recording heart rhythm was held to be processing health data because the recording could be used to predict arrhythmia.[127] The purposes for which consumer neurotechnology is marketed, sleep, meditation, mood management, cognitive enhancement, are similar to that of the wellbeing category, and the same reasoning suggests that neurodata collected for leisure, wellness, or lifestyle purposes ought to be treated as data concerning health and governed accordingly, notwithstanding the absence of any declared

[125] Medical Devices Regulation, (Regulation (EU) 2017/745)
[126] Kiesow Cortez (n 35).
[127] European Group on Ethics in Science and New Technologies (n 116).

medical purpose.[128] The same logic of function "by destination" exposes the corresponding weakness of intended-use classification under the Medical Devices Regulation, since a product framed as non-medical may in practice generate quasi-clinical inferences upon which users act, thereby approximating clinical assessment while escaping the framework designed to govern it.[129]

This boundary constituted a recurrent theme at the convening entitled "A European Strategy for Neurotechnology", hosted by the European Economic and Social Committee in Brussels on 13 April 2026 and organized together with the European Brain Council.[130] The scale of the wellness market was quantified upon that occasion. A study providing an analysis of close to three hundred dedicated neurotechnology companies worldwide, found that consumer firms now account for approximately sixty per cent of the sector and have more than tripled over the preceding decade, reaching market more rapidly than medical neurotechnology by reason of the smaller number of regulatory steps to which they are subject.[131] [132] Other studies show findings of greatest consequence for the gray zone, however, concerning disclosure: approximately half of the companies analyzed did not state within their terms and conditions that their products were not medically certified.[133] A competing concern was pressed with equal force at the same convening. It was observed that the existing governance instruments are already too numerous for founders to navigate, and cautioned against the creation of parallel tracks which would compound the perceived or real regulatory burden; the European Innovation Council representative noted that

[128] ibid.
[129] ibid.
[130] Centre for Future Generations & European Brain Council. (2026). A European Strategy for Neurotechnology [Convening, European Economic and Social Committee, Brussels, 13 April 2026]. Recording: https://www.youtube.com/watch?v=38E1m-5Gw6c
[131] CFG, 'Neurotech Consumer Market Atlas - How the Sector Is Making Moves into the Mainstream' (*Centre for Future Generations*, 16 June 2025) <https://cfg.eu/neurotech-market-atlas/> accessed 16 September 2026.
[132] European Brain Council (EBC) (n 115).
[133] ibid.

companies presently favour the pathway administered by the Food and Drug Administration by reason of the greater clarity of its contact points.[134] The response proposed was not the abandonment of oversight but rather a combination of measures, namely regulatory cohesion as between the medical and consumer device frameworks, regulatory sandboxes, earlier dialogue with regulators, common terminology, patient and clinician co-design, and voluntary soft-law mechanisms.[135] The soft-law mechanisms in question are principally the European Charter for the Responsible Development of Neurotechnologies[136], adopted by the European Brain Council in April 2025, and the Seal of Responsible Neurotechnology[137], launched at the convening as a means by which organisations might demonstrate alignment with the principles of the Charter in a verifiable manner.

4. European Health Data Space, research, and institutional governance

Regulation (EU) 2025/327 establishing the European Health Data Space applies from 26 March 2027 in staged form.[138] It is of consequence for present purposes because clinical neural data will increasingly be processed within a framework which specifies not merely legal permissions but also the architecture of data access for purposes of secondary use. The Regulation establishes governance bodies, cross-border infrastructures, health-data access mechanisms, and secure processing environments in respect of authorized secondary uses. Its importance to the argument advanced here is methodological rather than substantive: it demonstrates that European legislation

[134] ibid.
[135] CFG, 'Towards an EU Neurotechnology Strategy' (*Centre for Future Generations*, 10 April 2026) <https://cfg.eu/towards-an-eu-neurotechnology-strategy/> accessed 16 September 2026.
[136] European Brain Council (EBC), 'European Charter for the Responsible Development of Neurotechnologies' (23 September 2026) <https://www.braincouncil.eu/european-charter-for-the-responsible-development-of-neurotechnologies/> accessed 16 September 2026.
[137] Helena Kousi, 'EBC Introduces the Seal of Responsible Neurotechnology' (*European Brain Council (EBC)*, 9 April 2026) <https://www.braincouncil.eu/ebc-seal-of-responsible-neurotechnologies/> accessed 16 September 2026.
[138] (Regulation (EU) 2025/327)

is capable of moving beyond an abstract requirement of appropriate safeguards and of requiring instead that data be accessed and analyzed within a controlled technical environment rather than simply transmitted to each individual researcher.

In summary of the above discussion, the emerging European position is accordingly layered rather than neuro-specific in character. The General Data Protection Regulation governs the processing of personal data; the Artificial Intelligence Act restricts particular artificial intelligence uses; the Medical Devices Regulation governs qualifying devices together with Annex XVI products; the European Health Data Space structures the use of clinical health data; and fundamental-rights law[139] frames privacy, data protection, dignity, bodily and mental integrity, and freedom of thought. Institutional attention to the field has been extensive. The European Data Protection Supervisor has published dedicated technical analysis of neurodata[140]; the Council of Europe has assessed the impact of neurotechnology upon data protection[141]; and the research service of the European Parliament has examined the question of mental privacy specifically.[142] At the international level, the Organisation for Economic Co-operation and Development adopted the first intergovernmental instrument in the field[143], and the UNESCO Recommendation on the Ethics of Neurotechnology was adopted on 12 November 2025 at the forty-third session of the General Conference.[144]

[139] (Charter of Fundamental Rights & ECHR)
[140] 'TechDispatch #1/2024 - Neurodata | European Data Protection Supervisor' (n 2).
[141] Eduardo Bertoni and Marcello Ienca, 'The Privacy and Data Protection Implication of the Use of Neurotechnology and Neural Data from the Perspective of Convention 108' (Council of Europe 2024).
[142] European Parliament. Directorate General for Parliamentary Research Services., *The Protection of Mental Privacy in the Area of Neuroscience: Societal, Legal and Ethical Challenges.* (Publications Office 2024) <https://data.europa.eu/doi/10.2861/869928> accessed 30 September 2025.
[143] OECD (n 7).
[144] UNESCO (n 8).

The technical analysis of the European Data Protection Supervisor establishes several propositions upon which the present argument depends. It records that neurodata reliably permit the unique identification of individuals and are accordingly personal data, that they frequently constitute special categories of personal data whether as biometric data or as data concerning health, and that certain uses, among them lie detection and emotion inference in law-enforcement, workplace, and educational contexts, pose unacceptable risks to fundamental rights and are likely unlawful under Union law.[145] The same analysis draws attention to the difficulty which neurodata present for the principles of accuracy and of transparency, the former because the inferences drawn are probabilistic and the brain is subject to change over the life course, and the latter because the involuntary character of neural signals deprives the data subject of the direct control which the exercise of the rights under the General Data Protection Regulation presupposes.[146]

V. The American Framework

1. HIPAA: protection within a limited institutional perimeter

The Health Insurance Portability and Accountability Act of 1996 (HIPAA) is frequently invoked as though it constituted a general federal health-privacy statute. It does not. The Privacy, Security, and Breach Notification Rules promulgated under that Act apply to covered entities, being health plans, health care clearinghouses, and health care providers conducting covered electronic transactions, and to business associates performing specified functions or services on behalf of covered entities. Within that ecosystem, individually identifiable health information held or transmitted by the regulated actors may become protected health information, and electronic

[145] 'TechDispatch #1/2024 - Neurodata | European Data Protection Supervisor' (n 2).
[146] 'TechDispatch #1/2024 - Neurodata | European Data Protection Supervisor' (1 October 2025) <https:/
echDispatch #1/2024 - Neurodata | European Data Protection Supervisor' (n 2).

protected health information is subject in turn to administrative, physical, and technical safeguards. The framework additionally confers individual rights concerning access and imposes restrictions upon uses and disclosures.

Within the context of clinical neurotechnology the protection thus afforded may be substantial. An electroencephalographic recording, an implanted brain-computer interface recording, a neural diagnostic marker, or a decoded clinical output held by a hospital or by its business associate will typically fall within the framework governing protected health information. The boundary of that framework is, however, dependent upon the actor and upon the relationship rather than upon the sensitivity of the information. The Department of Health and Human Services has expressly explained that once information has been transmitted at the direction of an individual to an application which is neither a covered entity nor a business associate, the Rules generally cease to govern the downstream use made of that information by the application.[147] It follows that a consumer neurotechnology company does not become subject to regulation merely by reason of the circumstance that the information which it collects is physiologically or psychologically intimate.

That institutional perimeter gives rise to what may be described as the limitations of the US approach. A neural recording generated within a hospital may constitute protected health information, whereas an equivalent recording generated by a wellness headset may fall entirely outside the statutory scheme. The distinction rests not upon the inferential power of the neural signal but upon the identity and institutional role of the party collecting it. It is for this reason that the Health Insurance Portability and Accountability Act cannot serve as the sole federal answer to

[147] U.S. Department of Health and Human Services. (2025d). The access right, health apps, and APIs. https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/access-right-health-apps-apis/index.html

the question of mental privacy, notwithstanding that its concepts of security and of minimum necessary use remain instructive as design precedents. The Health Breach Notification Rule administered by the Federal Trade Commission, first adopted in 2009 and substantially amended in 2024, extends to qualifying vendors of personal health records and to related entities, and thereby narrows the gap in part in respect of applications and connected devices[148]; the Commission has in addition issued dedicated business guidance addressed to mobile health applications.[149] It remains, however, a breach-notification and consumer-protection instrument rather than a comprehensive code governing neural data.

2. FTC authority, FDA medical-device oversight, and the Common Rule

Outside the Health Insurance Portability and Accountability Act, the principal general authority of the Federal Trade Commission derives from Section 5 of the Federal Trade Commission Act, which prohibits unfair or deceptive acts or practices.[150] That authority is capable of addressing false privacy representations, undisclosed sharing, inadequate security, and certain unfair practices concerning health data. In the enforcement action brought against GoodRx, the Commission challenged disclosures of sensitive health information to advertising platforms and obtained undertakings concerning affirmative consent, deletion, retention, and the establishment of a privacy programme.[151] A consumer neurotechnology company which represented that it would maintain confidentiality while transmitting neural-derived health or mental-state information to

[148] 'Health Breach Notification Rule' (*Federal Trade Commission*, 19 September 2014) <https://www.ftc.gov/legal-library/browse/rules/health-breach-notification-rule> accessed 17 September 2026.
[149] 'Mobile Health App Interactive Tool' (*Federal Trade Commission*, 7 December 2022) <https://www.ftc.gov/business-guidance/resources/mobile-health-apps-interactive-tool> accessed 17 September 2026.
[150] (15 U.S.C. § 45).
[151] 'FTC Enforcement Action to Bar GoodRx from Sharing Consumers' Sensitive Health Info for Advertising' (*Federal Trade Commission*, 31 January 2023) <https://www.ftc.gov/news-events/news/press-releases/2023/02/ftc-enforcement-action-bar-goodrx-sharing-consumers-sensitive-health-info-advertising> accessed 17 September 2026.

advertising infrastructure would present a structurally similar problem. Section 5 is nevertheless less satisfactory as an instrument where a company proactively discloses expansive processing within dense contractual terms, since transparency does not of itself render a higher risk practice substantively acceptable.

The Food and Drug Administration regulates qualifying medical devices and has issued final guidance concerning implanted brain-computer interfaces intended for patients with paralysis or amputation.[152] That guidance addresses non-clinical testing and clinical-study considerations and reflects the circumstance that invasive brain-computer interfaces are presently moving from basic science towards clinical translation.[153] The oversight so exercised is, however, primarily a framework of product safety and efficacy; it does not establish a general privacy regime governing consumer neurotechnology.[154] Similarly, the Common Rule, codified at 45 C.F.R. Part 46, governs federally covered human-subject research through mechanisms of institutional review, informed consent, and specified exemptions, but it does not extend to every commercial research and development activity. The resulting arrangement is a mosaic in which medical, research, consumer, employment, and governmental contexts may trigger materially different legal obligations in respect of materially identical information.

[152] Center for Devices and Radiological Health, 'Implanted Brain-Computer Interface (BCI) Devices for Patients with Paralysis or Amputation - Non-Clinical Testing and Clinical Considerations' (17 May 2021) <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/implanted-brain-computer-interface-bci-devices-patients-paralysis-or-amputation-non-clinical-testing> accessed 17 September 2026.
[153] K Michelle Patrick-Krueger, Ian Burkhart and Jose L Contreras-Vidal, 'The State of Clinical Trials of Implantable Brain–Computer Interfaces' (2025) 3 Nature Reviews Bioengineering 50 <https://doi.org/10.1038/s44222-024-00239-5>.
[154] Wexler and Reiner (n 107).

3. The State-level neural-data laws

Colorado was the first state expressly to fold neural data into a comprehensive consumer-privacy statute. House Bill 24-1058, which took effect on 7 August 2024, amended the Colorado Privacy Act so that the expression "biological data" includes data generated by technological processing, measurement, or analysis of an individual's biological, genetic, biochemical, physiological, or neural properties, compositions, activities, body, or bodily functions, where such data is used or intended to be used for purposes of identification.[155] The statute separately defines neural data as information generated by measuring the activity of the central or peripheral nervous systems which is capable of being processed by or with the assistance of a device. The qualifier relating to identification purpose is decisive in its operation, since it is capable of narrowing statutory coverage even in circumstances where a neural signal supports sensitive affective or cognitive inference but has not been collected for purposes of identification.[156]

California's Senate Bill 1223, effective from 1 January 2025, amended the California Consumer Privacy Act, as modified by the California Privacy Rights Act, so as to include neural data within the category of sensitive personal information.[157] The California definition treats neural data as information generated by measuring the activity of a consumer's central or peripheral nervous system, and expressly excludes information inferred from nonneural information. Within the broader structure of the Act, the legal consequences which follow depend upon the manner in which sensitive personal information is collected and processed, including whether it is used for the purpose of inferring characteristics, and upon the consumer rights and controller duties thereby

[155] (Colorado H.B. 24-1058).
[156] 'The "Neural Data" Goldilocks Problem: Defining "Neural Data" in U.S. State Privacy Laws' (n 61).
[157] (California S.B. 1223).

triggered.[158] The source-based exclusion gives rise to a conceptual weakness of some importance: an emotional or cognitive state inferred from typing cadence, from eye movement, or from patterns of smartphone interaction may be functionally equivalent to an inference derived from electroencephalography while nevertheless falling outside the statutory definition of neural data.[159]

Montana's Senate Bill 163, effective from 1 October 2025, selected a different legislative home by proceeding to amend the state Genetic Information Privacy Act.[160] The statute defines neurotechnology and neurotechnology data and extends a consent structure modelled upon genetic privacy to information captured by neurotechnologies or generated from activity of the central or peripheral nervous system. Montana additionally defines nonneural information so as to include downstream physical effects of neural activity, which affords its boundary a greater degree of textual specificity than that achieved in California.[161] The statute reflects a deliberate analogy between genetic and neural information, both of which are biologically derived, potentially identifying, deeply inferential, and capable of revealing information not anticipated at the moment of collection.[162][163] It must be observed, however, that the analogy also imports the risk that regulation will direct its attention to the source of the information rather than to the consequences of the inference drawn from it.[164]

The amendments made to the Connecticut Data Privacy Act in 2025 took effect on 1 July 2026.[165] Connecticut defines neural data more narrowly than the states considered above, treating it as

[158] 'The "Neural Data" Goldilocks Problem: Defining "Neural Data" in U.S. State Privacy Laws' (n 61).
[159] Magee, Ienca and Farahany (n 88).
[160] (Montana S.B. 163).
[161] Clark, Linda, Crespo, Melissa and Wang, Katherine, 'Neural Data Added to Montana's Genetic Information Privacy Act' (*Morrison Foerster*, August 2025) <https://www.mofo.com/resources/insights/null> accessed 17 September 2026.
[162] Murray (n 69).
[163] Suter (n 90).
[164] Rothstein (n 65).
[165] (Connecticut Public Act 25-113).

information generated by measuring the activity of an individual's central nervous system, and expressly places such information within the category of sensitive data. The limitation to the central nervous system excludes peripheral measurement of the kind which Colorado, California, and Montana are each capable of capturing.[166] At the same time, the broader sensitive-data framework established in Connecticut now encompasses a range of health, biometric, financial, and identification information and requires consent in respect of the processing of sensitive data.[167] The four states therefore concur in according heightened treatment to neural data while diverging upon three matters of considerable practical consequence, namely the treatment of the central as against the peripheral nervous system, the treatment of inferred as against measured data, and the relevance of identification purpose.[168]

In June 2026, Vermont enacted Act 145 (S.71), defining neural data as information generated by measuring the activity of an individual's central nervous system and including such data within a comprehensive privacy framework, although the principal provisions of the enactment take effect on 1 January 2028.[169] Vermont additionally enacted Act 101 (H.814), which concerns neurological rights and the use of artificial intelligence within health and human services contexts.[170] Act 101 warrants a caveat which is easily overlooked. As introduced, the measure would have created enforceable neural-data privacy standards and would have prohibited devices from bypassing conscious decision-making without consent; those operative provisions were, however, removed in the course of passage. What was ultimately enacted is a declaratory recognition of rights to mental and neural data privacy and to freedom of thought, together with a directive to the state

[166] 'The "Neural Data" Goldilocks Problem: Defining "Neural Data" in U.S. State Privacy Laws' (n 61).
[167] (Connecticut Public Act 25-113).
[168] 'The "Neural Data" Goldilocks Problem: Defining "Neural Data" in U.S. State Privacy Laws' (n 61).
[169] (Vermont Act 145, 2026).
[170] (Vermont Act 101, 2026).

Artificial Intelligence Advisory Council to study the field and to report by 15 January 2027. It follows that Act 101 imposes no compliance obligations, and that it is Act 145 rather than Act 101 which alters the regulatory position in Vermont.

4. Federal proposals and constitutional dimensions

There exists at present no comprehensive federal statute governing the privacy of neural data. The Management of Individuals' Neural Data Act of 2025, introduced as S. 2925 on 29 September 2025, proposed that the Federal Trade Commission study the collection, processing, storage, transfer, and uses of neural and related data and report to Congress upon the gaps identified and the safeguards potentially available.[171] The scope of the proposal is notable in that it identifies advertising, employment, education, insurance, finance, behavioural influence, and commercial surveillance as distinct contexts rather than treating neurotechnology as a single undifferentiated market.[172] The proposal does not of itself supply a federal consent rule, a security standard, or a private right of action, but it is nevertheless indicative of the direction of federal concern. Professional bodies have moved in a parallel direction[173]

Constitutional doctrine assumes its greatest importance where government seeks access to neural information or seeks to compel a neural procedure. The analysis of self-incrimination offered by Farahany highlights the instability of the conventional distinction between physical evidence and testimonial communication in circumstances where a physical recording is capable of revealing

[171] Management of Individuals' Neural Data Act of 2025 (MIND Act 2025)
[172] 'Sens. Cantwell, Schumer, Markey Introduce Legislation to Shield Americans' Brain Data From Exploitation' (*U.S. Senate Committee on Commerce, Science, & Transportation*) <https://www.commerce.senate.gov/press/dem/release/sens-cantwell-schumer-markey-introduce-legislation-to-shield-americans-brain-data-from-exploitation/> accessed 17 September 2026.
[173] American Medical Association, 'Safeguarding neural data collected by neurotechnologies, Policy H-315.957, Resolution 503-A-25', 2025.

the content of recognition, memory, or intended speech.[174] The evidentiary question is not hypothetical, since neural measures have already been advanced as proof of memory,[175] and the issue was identified within the legal literature two decades ago.[176] Fourth Amendment doctrine becomes important within a distinct set of circumstances, namely where the state seeks to obtain an already-collected longitudinal record from a commercial provider.[177] These constitutional questions do not displace consumer privacy law. They demonstrate rather why commercial architecture is of consequence, since a constitutional right against compelled decoding is substantially weakened where the decoded output already exists within a vendor database and may be obtained by ordinary legal process.

5. Carpenter v. United States and revealing third-party records

Carpenter v. United States concerned the acquisition by government of historical cell-site location information held by wireless carriers. Federal investigators obtained one hundred and twenty-seven days of the defendant's location records by means of orders issued under the Stored Communications Act rather than by means of a warrant supported by probable cause.[178] The lower courts rejected the Fourth Amendment challenge upon the theory that location data was held by third parties and therefore fell within the conventional third-party doctrine. The Supreme Court granted review against a background of rapidly expanding digital records capable of reconstructing the movements of an individual over extended periods.

[174] Nita A Farahany, 'Incriminating Thoughts' (2012) 64 Stanford Law Review 351.
[175] Emily RD Murphy and Jesse Rissman, 'Evidence of Memory from Brain Data' (2020) 7 Journal of Law and the Biosciences lsaa078 <https://doi.org/10.1093/jlb/lsaa078>.
[176] 'Are Your Thoughts Your Own?:"Neuroprivacy" and the Legal Implications of Brain Imaging' (*New York City Bar Association*) <https://www.nycbar.org/reports/are-your-thoughts-your-ownneuroprivacy-and-the-legal-implications-of-brain-imaging/> accessed 17 September 2026.
[177] Farahany (n 6).
[178] (Carpenter v. United States, 585 U.S. 296 (2018)).

The Supreme Court held that the acquisition of the historical cell-site records constituted a search within the meaning of the Fourth Amendment and that the government was generally required to obtain a warrant. The majority emphasized the comprehensive, retrospective, and revealing character of long-term location records, and declined to extend the older third-party doctrine mechanically to the digital environment. It must be acknowledged that *Carpenter* did not concern neural data and that its holding should not to be overstated. Its relevance to the present analysis is analogical in character: continuous neural records may be at least as longitudinal and potentially more intimate than location history, and the circumstance that a vendor is in possession of such records should not automatically eliminate the constitutional significance of the information.[179] For purposes of governance design the more fundamental point is that minimization and local processing are capable of reducing the quantity of highly revealing data which subsequently becomes available to compulsory process.

6. Empirical audit evidence

The strongest argument in favour of operational standards is ultimately empirical rather than conceptual. The review of consumer neurotechnology companies conducted by the NeuroRights Foundation found widespread contractual rights of access, sharing, and reuse, inconsistent provision for deletion, and limited commitments to concrete technical safeguards.[180] Consumer devices are capable of collecting nervous-system data under contractual terms which permit substantial downstream use, and of doing so without offering the technical guarantees which would constrain that use.[181] This pattern is consistent with wider concerns expressed within the literature

[179] Farahany (n 6).
[180] Genser, Jared, Damianos, Stephen and Yuste, Rafael (n 45).
[181] ibid.

concerning direct-to-consumer neurotechnology[182], concerning neurosecurity and adversarial access to brain-computer interface systems[183] [184], and concerning the mismatch between notice-and-consent mechanisms and continuously generated physiological data.[185]

VI. Biometric and Genetic Analogies

Biometric law offers a structural analogy of some value, since biometric identifiers, like neural signals, may be persistent, difficult to alter, and useful both for purposes of identification and for the drawing of further inferences.[186] The Illinois Biometric Information Privacy Act requires notice and written consent prior to collection, mandates retention and destruction policies, and restricts sale and disclosure; unlike many consumer privacy statutes it additionally confers a private right of action.[187] The case law developed in Illinois illustrates the extent to which the practical force of the statute derives from that enforcement architecture rather than from its definitions. In *Rosenbach v. Six Flags* (2019)[188] the Supreme Court of Illinois held that a plaintiff need not plead actual injury beyond the statutory violation itself in order to maintain an action. In *Tims v. Black Horse Carriers* (2023)[189] the same court applied the five-year catch-all limitations period. In *Cothron v. White Castle* (2023)[190] it held that a separate claim accrues upon each scan, a

[182] Wexler and Reiner (n 5).
[183] Tamara Bonaci, Ryan Calo and Howard Jay Chizeck, 'App Stores for the Brain : Privacy and Security in Brain-Computer Interfaces' (2015) 34 IEEE Technology and Society Magazine 32 <https://doi.org/10.1109/MTS.2015.2425551>.
[184] Marcello Ienca and Pim Haselager, 'Hacking the Brain: Brain–Computer Interfacing Technology and the Ethics of Neurosecurity' (2016) 18 Ethics and Information Technology 117 <https://doi.org/10.1007/s10676-016-9398-9>.
[185] Viktor Ivanković, 'What (If Anything) Morally Separates Environmental from Neurochemical Behavioral Interventions?' (2023) 17 Neuroethics 6 <https://doi.org/10.1007/s12152-023-09540-3>.
[186] Quan, 'Neural Data Regulation: Lessons from Biometric Data's Evolution' (*California Lawyers Association*, 15 May 2026) <https://calawyers.org/privacy-law/neural-data-regulation-lessons-from-biometric-datas-evolution/> accessed 17 September 2026.
[187] (Illinois BIPA, 740 Ill. Comp. Stat. 14).
[188] Rosenbach v Six Flags Entertainment Corp, 2019 IL 123186, 129 NE3d 1197 (Ill 2019).
[189] Tims v Black Horse Carriers Inc, 2023 IL 127801, 216 NE3d 845 (Ill 2023).
[190] Cothron v White Castle System Inc, 2023 IL 128004, 216 NE3d 918 (Ill 2023).

construction of such financial consequence that the legislature amended the Act in the following year so as to limit recovery to a single violation per person per method of collection.[191] Federal standing doctrine has operated in the opposite direction, *Spokeo v. Robins* (2016)[192] and *TransUnion v. Ramirez* (2021)[193] having constrained which such claims may be litigated in federal court, and *Bryant v. Compass Group* (2020)[194] having sorted the notice-and-consent provisions of the Act from its retention-policy provision upon that basis. The lesson to be drawn is not that any one judicial interpretation is determinative. It is rather that enforcement architecture, retention rules, and the calibration of liability are capable of materially altering corporate behaviour, whereas a statutory definition standing alone does not specify the manner in which a lawful biometric template must technically be stored or computed upon.

Genetic privacy supplies the corresponding historical caution. Genetic exceptionalism produced special statutes because genetic information was understood to be unusually predictive, familial, permanent, and stigmatizing.[195] Critics subsequently argued that genetic-specific laws were liable to create a false sense of completion, to invite definitional disputes, and to leave functionally equivalent risks outside the protected category.[196] [197] Neural regulation faces the same danger in substantially the same form. A rule protecting only measurements of the nervous system is capable of failing to reach cognitive inferences drawn from eye movement, from typing, from speech, from facial behaviour, or from other data of the kind now described as cognitive biometrics.[198] It is

[191] (Public Act 103-0769).
[192] Spokeo Inc v Robins, 578 US 330 (2016).
[193] TransUnion LLC v Ramirez, 594 US 413 (2021).
[194] Bryant v Compass Group USA Inc, 958 F3d 617 (7th Cir 2020).
[195] Murray (n 69).
[196] Suter (n 90).
[197] Rothstein (n 65).
[198] Magee, Ienca and Farahany (n 88).

submitted that the stronger organizing principle is therefore inferential risk considered together with context, rather than neuroanatomical origin considered alone.[199]

A more recent institutional statement carries the same organizing principle further and in a direction which bears upon the framework proposed in Part IX below. The European Group on Ethics in Science and New Technologies has argued that the governance of neurotechnology cannot be confined to the individual device or to the isolated act of collection, because neural and behavioural data are increasingly aggregated into reusable infrastructures, large neural datasets and the Brain Foundation Models trained upon them, whose downstream inferences may affect persons and groups who never contributed any data at all.[200] Upon that view the appropriate object of governance shifts from the datum to the inference, and the Group accordingly proposes a principle of inferential proportionality, by which the permissibility of an increasingly powerful neurodata-derived inference is made to depend upon its scientific validity, its purpose, its context of use, and its foreseeable consequences, as a complement to, rather than a replacement for, the established principles of data minimization and purpose limitation.[201] This is the same conclusion reached by the analysis above from the direction of the genetic precedent: that the boundary of the category is the wrong axis, and that inferential risk considered together with context is the right one.

[199] Susser and Cabrera (n 28).
[200] European Group on Ethics in Science and New Technologies (n 116).
[201] ibid.

VII. Neural Decoding Capabilities

Every legal argument advanced within this field rests upon an empirical premise concerning capability, and there gaps in the legal literature in handling that premise, tending either to overclaim upon the strength of press coverage or to dismiss the field as speculative.

In order illuminate the status-quo of the neurotechnology capabilities, the following section focuses on what has presently been demonstrated, and proceeds in four parts addressing semantic decoding, inner speech, identifiability, and the resulting difficulty of categorization. This provides a basis for discussing further required technological safeguards per application capability which then can be implemented in legal standards in practice.

1. Semantic decoding and the cooperation constraint

The most frequently cited capability result is that of Tang et al. (2023)[202], who trained a decoder upon functional magnetic resonance imaging (fMRI) data capable of reconstructing intelligible word sequences which recover the meaning of perceived speech, of imagined speech, and of silent video. The headline finding is genuinely striking, since it concerns the recovery of continuous language by non-invasive means.

The finding of greatest legal consequence, however, is one which the authors deliberately tested and which the majority of legal citation omits. Recognizing that brain-computer interfaces should respect mental privacy, the authors examined whether decoding requires the cooperation of the subject, and found that it does, both for the purpose of training the decoder and for the purpose of applying it.[203] A subject who resists by performing a competing mental task is able to defeat the

[202] Tang and others (n 9).
[203] ibid.

decoder. This constitutes a technical defence of mental privacy which presently exists as a property of the method rather than as a designed safeguard, and the distinction between those two states of affairs is addressed later in this paper.

Défossez et al. (2023)[204] extended non-invasive decoding to magnetoencephalography and electroencephalography in respect of the perception of speech, achieving accuracies which permit identification of which sentence, among more than a thousand candidates, a subject is hearing. This constitutes a narrower capability than semantic reconstruction, but it has been demonstrated upon more accessible recording modalities, a circumstance of considerable regulatory importance given that consumer regulation must address wearable devices rather than research scanners.

Against these results there stands a methodological corrective, the scope of which requires precise statement. Shirakawa et al. examine[205] recent methods of text-guided visual image reconstruction which combine the Natural Scenes Dataset with text-to-image diffusion models, and demonstrate that their apparent success derives substantially from overlap between training and test categories and from generative hallucination rather than from genuine zero-shot decoding, a failure which the authors term output dimension collapse.[206] Their target is visual reconstruction rather than the language decoding reported by Tang et al. and by Défossez et al., and it would overstate their finding to treat it as a refutation of those results. What the critique does establish is a general methodological hazard applicable wherever a decoder is evaluated upon stimuli distributionally close to its training set, a condition which several neural decoding paradigms satisfy.

[204] Défossez and others (n 12).
[205] Ken Shirakawa and others, 'Spurious Reconstruction from Brain Activity' (2025) 190 Neural Networks 107515 <https://doi.org/10.1016/j.neunet.2025.107515>.
[206] ibid.

2. Inner speech and representational safeguards

An important result for the argument of this paper is on inner speech that is reported by Kunz et al. (2025).[207] Recording from motor cortex in participants with severe speech impairment, the authors found that inner speech is robustly represented and decodable in real time and, critically, that aspects of free-form inner speech could be decoded during sequence-recall and counting tasks which the participants had not intended to communicate. This constitutes genuine eavesdropping upon unintended content, and it accordingly moves the line established by Tang et al. (2023).

The same paper nevertheless reports a countervailing finding of equal importance. The authors identified a neural dimension corresponding to the intended overtness of speech, which distinguishes attempted communication from inner speech, and demonstrated that this dimension may be employed so as to prevent unintentional decoding.[208] They implemented what may be described as a passphrase gate, whereby the decoder activates only when the user thinks a specific and unusual phrase, which was recognised with high accuracy.

The importance of this result for present purposes is of high relevance in order to suggest a privacy safeguard to be implemented at the level of neural representation, that is to say, before the signal is captured for decoding rather than after it has been stored. Every mechanism considered in Part VIII operates downstream of capture: the data may be encrypted, perturbed, retained locally, or protected by the addition of noise to the model. The overtness gate operates upstream, constraining what the decoder is willing to read at all. This is relevant for legal scholarship as implementing the requirement of data protection by design would impact capabilities within this domain.

[207] Kunz and others (n 11).
[208] ibid.

3. Identifiability

Neural signals are reliably identifying, and the evidence supporting that proposition has become considerably stronger in recent years. Yang et al. (2022) recorded multi-task and multi-week electroencephalographic data[209] from ten pairs of monozygotic twins and isolated a person-identifying component which persisted across tasks and across weeks, constituting a structural brainprint rather than a task-dependent one, and found that twin siblings could nevertheless be distinguished from one another, most effectively by reference to signal originating in late-developing cortical areas.[210] Wu, Ramdas and Wehbe (2022) demonstrate a comparable result[211] in respect of magnetoencephalography, showing that simple interpretable spatial, temporal, and frequency features identify individuals across days, across tasks, and even across recording modality as between magnetoencephalography and electroencephalography. They frame that finding explicitly as a privacy risk arising from the sharing of nominally anonymised brain data.[212] An entire subfield treats brainprints as an attractive biometric precisely due to the difficulty of spoofing them.[213]

Colorado regulates biological data used for purposes of identification.[214] If neural signals are inherently identifying, the qualifier either performs no work, since every neural recording

[209] Yao-Yuan Yang and others, 'Person-Identifying Brainprints Are Stably Embedded in EEG Mindprints' (2022) 12 Scientific Reports 17031 <https://doi.org/10.1038/s41598-022-21384-0>.
[210] ibid.
[211] Shenghao Wu, Aaditya Ramdas and Leila Wehbe, 'Brainprints: Identifying Individuals from Magnetoencephalograms' (2022) 5 Communications Biology 852 <https://doi.org/10.1038/s42003-022-03727-9>.
[212] ibid.
[213] Xia and others (n 63).
[214] (Colorado H.B. 24-1058).

potentially serves identification, or else it performs too much, in that it permits a controller to disclaim identification purpose while holding data which in fact identifies.[215] [216]

Identifiability additionally disposes of anonymisation as a route to compliance. A controller asserting that neural recordings have been anonymised is advancing a claim of the type which Ohm (2010) demonstrated to be generally unreliable[217], and which Wu et al. (2022) demonstrate[218] concretely in respect of shared neuroimaging datasets.

4. The category problem

Magee, Ienca, and Farahany (2024) advance the argument[219] which has implications for the drafters of neural data statutes. They propose cognitive biometrics as the operative category, upon the ground that data derived from wearables, from extended-reality systems, and from interaction telemetry, encompassing eye movement, pupil dilation, keystroke dynamics, and heart rate variability, supports inferences concerning cognitive, affective, and conative states which are comparable to those derived from neural measurement.

If that proposition is correct, it follows that statutes drawn around the measurement of the nervous system have drawn the boundary in the wrong place. The express exclusion in California of information inferred from nonneural sources exempts precisely the category which carries a substantial proportion of the risk[220] The limitation in Connecticut to the central nervous system operates to the same effect.[221] The use-based framing adopted in the Artificial Intelligence Act,

[215] 'The "Neural Data" Goldilocks Problem: Defining "Neural Data" in U.S. State Privacy Laws' (n 61).
[216] Kobbi Nissim and others, 'Bridging the Gap between Computer Science and Legal Approaches to Privacy' (2017) 31 Harvard Journal of Law & Technology (Harvard JOLT) 687.
[217] Paul Ohm, 'Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization' (2010) 57 UCLA Law Review 1701.
[218] Wu, Ramdas and Wehbe (n 211).
[219] Magee, Ienca and Farahany (n 88).
[220] (California S.B. 1223).
[221] (Connecticut Public Act 25-113).

which prohibits emotion inference irrespective of the source of the signal, is by this measure better drafted than the American source-based provisions.[222] The same point was pressed from the data-protection side at the April 2026 European convening and it was argued that the gap within the General Data Protection Regulation concerns inferences about the mind rather than neurotechnology as such.[223]

The observation reinforces the central claim of this paper. If the boundary of the category cannot reliably be drawn around the source of the data, then protection must attach instead to what is done with the data and to the manner in which the processing is technically constructed, which is to say to inferential purpose and to architecture rather than to origin.

VIII. Privacy-Enhancing Technologies and Their Legal Significance

1. Cryptographic approaches

Homomorphic encryption permits computation to be conducted upon encrypted data without decryption. Popescu et al. (2021) demonstrated its application to electroencephalography[224], developing an encoding scheme which permits standard homomorphic encryption to operate upon real-valued data of arbitrary precision, and evaluating that scheme upon tasks of seizure detection and of prediction of alcoholism predisposition. Their reporting of accuracy and of computational cost is objective, since a legal standard mandating encryption in use rather than merely encryption at rest imposes real performance costs which a proportionality analysis might weigh.

[222] (Regulation (EU) 2024/1689).
[223] European Brain Council (EBC) (n 115).
[224] Andreea Bianca Popescu and others, 'Privacy Preserving Classification of EEG Data Using Machine Learning and Homomorphic Encryption' (2021) 11 Applied Sciences 7360 <https://doi.org/10.3390/app11167360>.

Secure multiparty computation distributes computation across parties in such a manner that no party learns the inputs of any other. Agarwal et al. (2019) applied secret-sharing protocols to brain-computer interface regression[225], thereby permitting a service provider to run a model upon the neural signals of a user without observing those signals, and without disclosing the proprietary model to the user. Their stated motivation is the leakage of credentials and of other private information inferable from neural signals. The result establishes that obligations of data minimization would not require the abandonment of cloud analytics, which is a conclusion of considerable practical significance.

2. Differential privacy and synthetic data

Differential privacy provides a formal and composable guarantee bounding what any output may reveal concerning any individual record.[226] Applications of the technique to neural data have concentrated upon the generation of synthetic data and upon federated training.[227] [228] Paul & Bajwa (2026)[229] combine differentially private stochastic gradient descent with a generative adversarial network in order to synthesize electroencephalographic data, evaluate the result, and report improved robustness to membership inference and reconstruction attacks by comparison with earlier approaches. The design point which carries over to the legal analysis is that their privacy

[225] Anisha Agarwal and others, 'Protecting Privacy of Users in Brain-Computer Interface Applications' (2019) 27 IEEE Transactions on Neural Systems and Rehabilitation Engineering 1546 <https://doi.org/10.1109/TNSRE.2019.2926965>.
[226] Cynthia Dwork, 'Differential Privacy: A Survey of Results' in Manindra Agrawal and others (eds), *Theory and Applications of Models of Computation* (Springer 2008) <https://doi.org/10.1007/978-3-540-79228-4_1>.
[227] H Brendan McMahan and others, 'Communication-Efficient Learning of Deep Networks from Decentralized Data' (2017) Proceedings of the 20 th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017. JMLR: W&CP volume 54 <https://research.google/pubs/communication-efficient-learning-of-deep-networks-from-decentralized-data/> accessed 17 September 2026.
[228] Xia and others (n 63).
[229] Shouvik Paul and Garima Bajwa, 'Q-DP-GAN: Improving EEG Data Privacy Through Quantum-Inspired Differential Privacy-Based GAN' (2026) 10 Cryptography 31 <https://doi.org/10.3390/cryptography10030031>.

parameters are adjusted dynamically by means of adaptive noise and clipping rather than being fixed in advance.

The recurring finding across this literature is that the privacy budget is most usefully treated as adaptive rather than as fixed, because the utility cost of a given level of noise varies substantially with the task and with signal quality.[230] This could have an implication for legislative drafting, i.e. a statute specifying a numerical value would be too extensive, however the law may sensibly require a process, i.e. a budget be set, justified, documented, and periodically reviewed.[231]

3. Local processing and perturbation

Edge processing addresses the problem at its source, by extracting features upon or in proximity to the device so that raw signal never transits to a server. Yan et al. (2023)[232] propose a scheme designated LightPyFE, being a lightweight privacy-preserving feature-extraction scheme for electroencephalography under edge computing. An outsourced computing toolkit constructed upon additive secret sharing and secure interaction protocols permits edge servers to perform secure integer and floating-point operations while the user retains encryption and decryption, with the consequence that the servers never observe raw electroencephalographic data. The scheme addresses latency and privacy jointly, which is the circumstance rendering it the strongest available technical basis for a legal default of local processing. Two limitations require to be noted are as follows: the scheme covers feature extraction only and does not extend to encrypted classification; and it assumes a threat model in which edge servers do not actively attempt inference. A legal

[230] ibid.
[231] Information Technology Laboratory Computer Security Division, 'Privacy-Enhancing Cryptography | CSRC | CSRC' (*CSRC | NIST*, 3 January 2017) <https://csrc.nist.gov/projects/pec> accessed 17 September 2026.
[232] Nazhao Yan and others, 'Lightweight Privacy-Preserving Feature Extraction for EEG Signals Under Edge Computing' (2024) 11 IEEE Internet of Things Journal 2520 <https://doi.org/10.1109/JIOT.2023.3292232>.

presumption of locality constructed upon this evidence ought accordingly to be framed as a presumption concerning the locations to which raw signal may travel, and not as a claim that any particular scheme discharges the obligation in full.

Perturbation approaches proceed differently, modifying the signal so as to suppress specific attributes while preserving utility in respect of the primary task. Meng et al. (2024) first demonstrate[233] empirically that user identity, gender, and prior experience of brain-computer interfaces are each inferable from electroencephalography recorded for an unrelated task, and then design perturbations which reduce those inferences to near chance while preserving accuracy upon the primary task. The demonstrative half of that contribution is as valuable to lawyers as the mitigating half, since it refutes the claim that task-limited collection is inherently protective of privacy.

4. The need for a synthesis approach

Federated learning, which is frequently invoked as a solution upon the ground that raw data never leaves the device[234], does not of itself provide a privacy guarantee. Membership inference attacks determine whether a particular record participated in training and operate against federated models both during training and following deployment; Nasr et al. (2019) demonstrate[235] this systematically within federated settings by exploiting per-sample gradients, intermediate

[233] Lubin Meng and others, 'Protecting Multiple Types of Privacy Simultaneously in EEG-Based Brain-Computer Interfaces' (arXiv, 29 November 2024) <https://doi.org/10.48550/arXiv.2411.19498> accessed 17 September 2026.
[234] McMahan and others (n 227).
[235] Milad Nasr, Reza Shokri and Amir Houmansadr, 'Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-Box Inference Attacks against Centralized and Federated Learning' *2019 IEEE Symposium on Security and Privacy (SP)* (2019) <https://doi.org/10.1109/SP.2019.00065> accessed 17 September 2026.

activations, and loss values observable to a participating client or to the aggregating server. Gradient inversion attacks reconstruct training inputs from model updates.[236]

It must be acknowledged that quantified attack results in respect of federated neural data specifically have not yet been published. The attacks described above have been evaluated upon image, text, and tabular benchmarks, and their transfer to the signal characteristics, client counts, and update schedules typical of neural datasets is an assumption rather than a finding.

The regulatory implication is that any compliance regime which certifies a single mechanism will fail. Brauneck et al. (2023), reviewing federated learning and privacy-enhancing technologies against the requirements of data protection law in medical research[237], reach the constructive version of this conclusion, namely that combinations of federated learning with secure multiparty computation and differential privacy are capable of providing safeguards sufficient to satisfy data protection requirements at comparable levels of model performance. Combination rather than selection is therefore the operative concept, and it must be added that any such combination requires to be revisable as attacks improve.

The broader privacy-engineering literature reinforces the proposition that privacy-enhancing technologies ought to be treated as properties and processes and not solely focused on as labels. The Advanced Encryption Standard and Transport Layer Security supply recognized

[236] Model inversion attacks, established within the general machine-learning literature by Fredrikson et al. (2015), recover information concerning training data from the outputs of a deployed model. Matt Fredrikson, Somesh Jha and Thomas Ristenpart, ‘Model Inversion Attacks That Exploit Confidence Information and Basic Countermeasures’ Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (Association for Computing Machinery 2015).

[237] Alissa Brauneck and others, ‘Federated Machine Learning, Privacy-Enhancing Technologies, and Data Protection Laws in Medical Research: Scoping Review’ (2023) 25 Journal of Medical Internet Research e41588 <https://doi.org/10.2196/41588>.

confidentiality primitives in respect of data at rest and in transit.[238] Differential privacy is capable of bounding information leakage from statistics or from model training where its privacy unit, adjacency relation, and cumulative budget are correctly specified.[239] [240] Federated learning is capable of retaining raw data locally but may nevertheless leak through gradients or updates.[241] Secure aggregation is capable of concealing participant-level updates.[242] Homomorphic encryption and secure multiparty computation are capable of permitting analysis without general plaintext access[243][244] None of these techniques, it must be emphasized, renders an otherwise unlawful purpose lawful therefore the focus in this section is on the security of legitimately collected data.

Within the context of neural systems, local and on-device processing deserves particular weight by reason of its capacity to eliminate entire data flows rather than merely to protect them.[245] Where a headset is capable of converting raw electroencephalographic signal into a cursor command or an attention score locally, the transmission of only the narrow functional output reduces exposure to breach, to insider access, to secondary commercialization, and to compulsory-process risk.[246] Signal perturbation is capable of suppressing identity and demographic attributes while preserving a primary task,[247] and representation-level gates are capable of constraining a decoder to intended

[238] National Institute of Standards and Technology, 'Advanced Encryption Standard (AES)' (US Department of Commerce 2023) Federal Information Processing Standard (FIPS) 197 <https://doi.org/10.6028/NIST.FIPS.197-upd1> accessed 17 September 2026.
[239] Dwork (n 226).
[240] Joseph Near, 'Protecting Model Updates in Privacy-Preserving Federated Learning' [2024] NIST <https://www.nist.gov/blogs/cybersecurity-insights/protecting-model-updates-privacy-preserving-federated-learning> accessed 17 September 2026.
[241] McMahan and others (n 227).
[242] Near (n 241).
[243] Popescu and others (n 224).
[244] Agarwal and others (n 225).
[245] Yuste and others (n 39).
[246] Yan and others (n 232).
[247] Meng and others (n 233).

communication.[248] These techniques operationalize data minimization more robustly than do contractual prohibitions applied after raw data has already been centralized.[249]

IX. Translating Legal Standards into Technical Requirements

1. The formalization method

The methodological resource which this field requires already exists, having been developed within an adjacent domain. Nissim et al. (2018) set out a two-step method for bridging computer-science and legal approaches to privacy.[250] The first step consists in extracting from a legal standard a formal mathematical requirement, adopting a conservative worst-case reading where the legal text is ambiguous. The second consists in constructing a rigorous proof that a proposed technical mechanism satisfies that requirement. The authors demonstrated the method upon the Family Educational Rights and Privacy Act and differential privacy.

Cohen and Nissim (2020) demonstrated that the method is logically sound.[251] They formalized the notion of singling out employed within the General Data Protection Regulation as predicate singling out, proved that any release claimed to be anonymous must be secure against it, and demonstrated that differential privacy provides such security whereas k-anonymity does not, thereby approaching differently to the position which had been taken by the Article 29 Working Party.[252]

[248] Kunz and others (n 11).
[249] Xia and others (n 63).
[250] Nissim and others (n 216).
[251] Cohen and Nissim (n 102).
[252] It should be noted that the referred work of the Working Party dates back to 2014, therefore it is probable that the advancement of the techniques might not have been foreseen at the time. Article 29 Data Protection Working Party. (2014). Opinion 05/2014 on Anonymisation Techniques (WP216, 0829/14/EN). European Commission.

This might constitute a promising model in respect of neural data. The relevant legal standards, the requirement of security appropriate to the risk under Article 32, the state of the art under Article 25, the reasonable security required by the American statutes, and the identification purpose qualifier employed in Colorado, are precisely the kind of underspecified standards which the method was designed to address.[253] [254]

2. Standards & Specification

The critique of Article 25 offered by Waldman (2020)[255] identifies the mechanism of failure with precision: a standard sufficiently vague to be adaptable is also sufficiently vague to be absorbed into compliance practices without producing substantive effect. The guidance issued by the European Data Protection Board upon Article 25 elaborates the concept of state of the art as a moving benchmark referenced to available and effective measures, which is the correct structure but which requires an authoritative account of what is available and effective within a given domain.[256] A comparable critique, accompanied by a constructive proposal for reform, has been offered elsewhere within the literature.[257] In respect of neural data no such account presently exists in any form which a supervisory authority could apply, although it should be noted that the pseudonymization guidelines issued by the Board in 2025 demonstrate that it is willing to give technique-level content to an Article 32 measure where it elects to do so.[258]

The literature concerning anonymization supplies the corrective necessary to any proposal that a single technique be mandated. Ohm (2010) established that anonymization fails against

[253] (Regulation (EU) 2016/679).
[254] (Colorado H.B. 24-1058).
[255] Waldman (n 62).
[256] European Data Protection Board, 'Guidelines 4/2019 on Article 25 Data Protection by Design and by Default' (n 60).
[257] Rubinstein and Good (n 101).
[258] European Data Protection Board, 'Guidelines 01/2025 on Pseudonymisation' (n 99).

determined re-identification.[259] Rubinstein and Hartzog (2016) drew the constructive conclusion that protection ought to be reconceived as process-based risk management rather than as a binary release decision.[260] Applied to the present context, the conclusion is that the law ought to require a defensible process for selecting, combining, and revising technical measures against a documented threat model, rather than certifying any particular mechanism.

3. A graded framework

Synthesizing the foregoing contributions, five design principles may be articulated. The first four follow from the literature considered above; the fifth is novel to the argument of this paper.

First, the strength of protection ought to scale with inferential risk rather than with recording modality. What matters is not whether a signal originates within the central nervous system but what may be inferred from it, by whom, and with what consequence. This principle answers the definitional problem identified[261] and avoids the boundary difficulties encountered in Colorado and in Connecticut.[262]

This first principle coincides with the principle of inferential proportionality advanced by the European Group on Ethics in Science and New Technologies, which likewise directs the strength of protection to the inference drawn and its consequences rather than to the modality of recording, and which extends the analysis beyond the individual data subject to the reuse of neurodata within foundation models and shared infrastructures whose effects fall upon others.[263]

[259] Ohm (n 217).
[260] Ira S Rubinstein and Woodrow Hartzog, 'Anonymization and Risk' (2016) 91 Washington Law Review 703.
[261] Magee, Ienca and Farahany (n 88).
[262] (Colorado H.B. 24-1058) & (Connecticut Public Act 25-113).
[263] European Group on Ethics in Science and New Technologies (n 116).

Second, processing locality ought to constitute the default position, with transmission requiring justification. Yuste et al. (2017) establish that this choice is architecturally consequential[264], and it might be difficult to reverse once made. The technical feasibility of the approach is demonstrated by edge feature-extraction schemes of the kind proposed.[265] Where transmission is necessary, as it will be in respect of the training of foundation models[266], the burden should fall upon the controller to justify it and to specify compensating measures.

Third, mechanisms ought to be combined, and the combination selected should be justified against an explicit threat model[267] and a single technique might not be capable of discharging the obligation in full.[268]

Fourth, obligations should be procedural and revisable rather than fixed. A mental data protection impact assessment[269] constitutes the natural vehicle, since it requires the controller to document the threat model, the mechanisms selected, the residual risk, and the schedule of review, and since it produces a record which a supervisory authority is able to audit.[270]

Fifth, and as the result reported by Kunz et al. (2025) indicates[271], where the technology permits, safeguards should be implemented at the point of signal interpretation rather than solely at the point of storage or transmission. An overtness gate constitutes a stronger protection than any

[264] Yuste and others (n 39).
[265] Yan and others (n 232).
[266] Wei-Bang Jiang, Li-Ming Zhao and Bao-Liang Lu, 'Large Brain Model for Learning Generic Representations with Tremendous EEG Data in BCI' (arXiv, 29 May 2024) <https://doi.org/10.48550/arXiv.2405.18765> accessed 17 September 2026.
[267] Brauneck and others (n 238).
[268] Nasr, Shokri and Houmansadr (n 236).
[269] Ienca and Malgieri (n 93).
[270] (Regulation (EU) 2016/679).
[271] Kunz and others (n 11).

downstream mechanism, for the straightforward reason that data never captured cannot be breached, subpoenaed, or repurposed.

A workable compliance model ought in addition to distinguish baseline controls from risk-scaled controls. Baseline controls in respect of identifiable neural information should include encryption in transit and at rest, key separation and rotation, access logging, deletion and backup governance, versioned documentation of data flows, and testing following material software changes. Higher-risk uses ought to add a presumption of locality, secure aggregation or encrypted computation in respect of collaborative analysis,[272] [273] differential privacy in respect of population-level releases where suitable,[274] and controls directed against unauthorized secondary inference.[275] Rather than writing a particular cipher or epsilon value into legislation, regulators ought to maintain updateable technical profiles and to require controllers to justify deviations within a neural-data protection impact assessment.[276][277]

A controller which follows a recognized profile ought to gain a rebuttable presumption in respect only of the specific duty which the profile tests, confidentiality or minimization, for example, rather than a general safe harbour in respect of the processing activity considered as a whole. The presumption ought to disappear where implementation is defective, where parameters are unjustified, where undeclared telemetry exists, or where the controller repurposes data for an unlawful or prohibited inference. This approach converts privacy engineering into a compliance

[272] Near (n 241).
[273] Agarwal and others (n 225).
[274] Joseph Near and others, 'Guidelines for Evaluating Differential Privacy Guarantees' (National Institute of Standards and Technology 2025) NIST Special Publication (SP) 800-226 <https://doi.org/10.6028/NIST.SP.800-226> accessed 16 September 2026.
[275] Meng and others (n 233).
[276] (Regulation (EU) 2016/679).
[277] European Data Protection Board, 'Guidelines 4/2019 on Article 25 Data Protection by Design and by Default' (n 60).

incentive while preserving substantive limits upon advertising, discrimination, manipulation, and coercive monitoring. The underlying economics, namely that data is nonrival and that privacy involves genuine welfare trade-offs rather than pure loss, supports calibration rather than prohibition.[278] [279]

An instructive instance of the mechanism proposed here is already under construction, although it has not been presented in these terms. The Seal of Responsible Neurotechnology pilot operated by the European Brain Council translates the nine ethical areas of its Charter into structured self-assessment checklists which are designed for the use of the organizations who are developing medical, non-medical and research neurotechnologies.[280] It is designed to serve as an ethical label confirming that the organization is aligned with the values of the Charter, willing to commit to a voluntary assessment. Corresponding with broader data regulations in the EU, the criteria for the Seal of Responsible Neurotechnology is intended to cover areas such as purpose, benefit and proportionality of data collection, neural data governance and cybersecurity, prevention of misuse and accountability.[281] Regarding cybersecurity the criteria includes "treatment of neural and neural-adjacent data as sensitive personal data, with commensurate safeguards on collection, storage, processing, and secondary use"[282] therefore further enhancing the need for the work on a technical framework that would be supported by legal instruments. This initiative is intended to pilot in 2026 and operate at two tiers, the Seal of Trust being awarded upon satisfaction of the essential criteria and a Seal of Excellence being reserved for organizations exceeding them, hence

[278] Charles I Jones and Christopher Tonetti, 'Nonrivalry and the Economics of Data' (2020) 110 American Economic Review 2819 <https://doi.org/10.1257/aer.20191330>.
[279] Alessandro Acquisti, Curtis Taylor and Liad Wagman, 'The Economics of Privacy' (2016) 54 Journal of Economic Literature 442 <https://doi.org/10.1257/jel.54.2.442>.
[280] European Brain Council. (2026) (n 130).
[281] ibid.
[282] ibid.

creating further incentives for organizations to demonstrate a commitment to responsible technology as well as to regulatory foresight.[283]

X. Sector-Specific Application

This Part applies the framework developed above to five sectors, namely clinical neurotechnology, employment, consumer devices, research, and education. The sectors are treated separately because the consent mechanism operates with materially different significance within each of them, and because the appropriate default architecture varies correspondingly.

1. Clinical neurotechnology

Clinical applications present the least difficult case. Data is generated within a regulated relationship, the purpose of processing is documented, and within Europe both the medical devices regime and, in respect of secondary use, the European Health Data Space are applicable.[284] The secure processing environment established under the European Health Data Space already instantiates the principle of locality in respect of research use.

The residual issue concerns secondary use, and specifically the application of Article 89(2) to commercially funded research conducted upon clinical neural data, a question which was flagged[285] and which remains unresolved. It may be noted that the draft guidelines issued by the European Data Protection Board in 2026 concerning processing for scientific research purposes constitute the first authoritative text likely to bear upon that question.[286] Given the identifiability

[283] ibid.
[284] (Regulation (EU) 2017/745) & Regulation (EU) 2025/327).
[285] Ienca and others (n 55).
[286] European Data Protection Board, 'Guidelines 1/2026 on Processing of Personal Data for Scientific Research Purposes | European Data Protection Board' (25 June 2026) <https://www.edpb.europa.eu/public-consultations/guidelines-12026-on-processing-of-personal-data-for-scientific-research_en> accessed 17 September 2026.

of neural signals established in Part VII above, a research derogation premised upon anonymization is not robust.[287] The appropriate requirement is accordingly architectural in character, consisting in secure processing environments with output checking, combined with differential privacy in respect of any released aggregate.[288]

Within the United States, the Health Insurance Portability and Accountability Act governs the clinical setting adequately at the level of disclosure control but is silent as to processing architecture, and its de-identification safe harbour is, in respect of neural signals, the same unsound proposition.

2. Employment

The workplace is the sector in which the case for stringent requirements is strongest, for the reason that the consent mechanism fails there most completely. Muhl (2024) analyses wearable neurodevices employed for workplace monitoring under both the General Data Protection Regulation and the Artificial Intelligence Act, concluding that European law affords substantial protection while noting the attendant risk of over-restriction in respect of legitimate safety applications.[289] The power asymmetry inherent within the employment relationship means that consent is rarely freely given, and it is for that reason that Article 5(1)(f) of the Artificial Intelligence Act adopts a prohibition rather than a consent requirement when it comes to emotion inference.[290]

The prohibition does not, however, extend to neural devices employed for purposes other than emotion inference, such as fatigue detection in safety-critical roles or device control. In respect of

[287] Yang and others (n 209).
[288] Near and others (n 275).
[289] Muhl (n 112).
[290] (Regulation (EU) 2024/1689).

such uses the framework advanced here prescribes the most stringent settings available, namely on-device processing with no transmission of raw signal,[291] extraction limited to the specific operational feature, perturbation of the transmitted feature so as to suppress identity and demographic attributes,[292] and prohibition of any secondary use.

Within the United States, where no equivalent prohibition presently exists, the employment sector constitutes the strongest candidate for a use-based statutory rule. The biometric experience is directly relevant in this connection, since the most consequential litigation under the Illinois statute concerned workplace fingerprint timekeeping,[293] and since workplace deployment is the context in which neural monitoring is likely first to become routine.[294]

3. Consumer devices

Consumer neurotechnology constitutes the largest category by volume, the least regulated, and the category in respect of which the available audit evidence discloses the worst practice.[295] It is additionally the fastest-growing, consumer firms now accounting for approximately sixty per cent of the sector.[296] It is the category in which the four American definitions diverge most consequentially[297] and in which European protection rests upon the least determinate provisions.

The prescription advanced here consists in a graded default connected to inferential purpose. A device performing device control requires on-device processing and the transmission of control signals only. A device performing state classification, whether of attention, of focus, or of sleep

[291] Yan and others (n 232).
[292] Meng and others (n 233).
[293] Cothron v White Castle System Inc, 2023 IL 128004, 216 NE3d 918 (Ill 2023).
[294] Muhl (n 112).
[295] Genser, Jared, Damianos, Stephen and Yuste, Rafael (n 45).
[296] European Brain Council (EBC) (n 115).
[297] 'The "Neural Data" Goldilocks Problem: Defining "Neural Data" in U.S. State Privacy Laws' (n 61).

stage, requires on-device classification with transmission of the classification rather than of the underlying signal, together with a documented justification in respect of any deviation.[298] A device the purpose of which is inference concerning affective or cognitive traits ought to be presumptively prohibited outside clinical settings, following the model adopted within the Artificial Intelligence Act[299] but extended beyond the contexts of workplace and education.

The question of enforcement is unavoidable. Upon the biometric evidence set out in Part VI above, obligations unaccompanied by a private right of action or by credible public enforcement produce limited alteration in behaviour.[300] Given the Illinois experience with the calibration of liability, and in particular the per-scan accrual rule adopted in case law[301] together with the rapid legislative correction of that rule effected by Public Act 103-0769, the appropriate design consists in a private right of action with damages calculated per violation rather than per measurement, coupled with a safe harbour available to controllers able to demonstrate compliance with a specified technical standard.

4. Research

Research data sharing poses challenges to the framework considered as a whole. Neural foundation models require large pooled cross-subject corpora. Jiang, Zhao, and Lu (2024) pre-trained their Large Brain Model[302] upon approximately two thousand five hundred hours of heterogeneous electroencephalographic data drawn from around twenty datasets, and report that it outperforms task-specific state-of-the-art models across abnormality detection, event classification, emotion

[298] Yan and others (n 232).
[299] (Regulation (EU) 2024/1689).
[300] Rosenbach v Six Flags Entertainment Corp, 2019 IL 123186, 129 NE3d 1197 (Ill 2019).
[301] Cothron v White Castle System Inc, 2023 IL 128004, 216 NE3d 918 (Ill 2023).
[302] Jiang, Zhao and Lu (n 267).

recognition, and gait prediction, which could be viewed as a default of local processing standing in conflict with that trajectory.

This constitutes the strongest objection to the framework advanced here, and it is an objection originating upon the technical side rather than in industry advocacy, a circumstance which renders it correspondingly more difficult to dismiss. Three responses are available. First, federated and split-learning architectures permit pooled training without pooled raw data,[303] albeit at some cost in performance and with the residual attack surface documented in Part VIII above. Second, differentially private synthetic data permits the construction of a corpus subject to a formal bound upon individual disclosure.[304] Third, the secure processing environment model adopted under the European Health Data Space[305] permits centralized analysis under architectural constraint rather than by means of raw-data transfer.

Although these responses do not resolve the tension entirely, a default of locality imposes a real cost upon model development, that the magnitude of that cost is declining as federated methods improve, and that the appropriate regulatory posture consists in requiring justification and compensating measures with regard to pooling rather than in prohibiting it.

5. Education

Education merits separate treatment for two reasons: consent is structurally weak within the educational relationship, and scores relating to attention or emotion are capable of influencing consequential judgments concerning children. Consent of the consumer variety cannot cure a school mandate requiring the wearing of an attention-monitoring headset. The BrainCo example

[303] McMahan and others (n 227).
[304] Paul and Bajwa (n 229).
[305] (Regulation (EU) 2025/327).

illustrates why a technically modest electroencephalographic classification may become socially intrusive once it is embedded within a hierarchy and displayed as a measure of attentiveness.[306] The Focus headband manufactured by that company, which scores attention from electroencephalographic signal and reports the resulting score to a teacher-facing dashboard, was deployed within classrooms in Zhejiang province; following the circulation of photographs in 2019 the local education bureau suspended its use pending review of the manner in which the collected data was being handled. It is worth emphasizing that the technical claim at issue was modest, consisting in a coarse attention index rather than in decoded content, and that this is precisely the point: the harm lay in the display of the score within a relationship of authority rather than in the inferential power of the signal itself.[307]

Within Europe, the prohibition contained within the Artificial Intelligence Act upon specified emotion recognition in educational institutions supplies a use-based constraint; within the United States, protection depends upon a combination of state privacy law, education law, contract, and institutional policy. A sectoral rule ought accordingly to prohibit compulsory neural monitoring for purposes of discipline, of grading, or of generalized assessment of attentiveness, subject only to narrow exceptions relating to accessibility, to medical need, or to safety, and subject in each case to independent review, and recently federal attention to the educational context is already discernible within the sectors enumerated in the MIND Act.

XI. Towards Transatlantic Harmonization

[306] 'Chinese School Forced to Stop Using Headbands to Monitor Concentration' (*South China Morning Post*, 1 November 2019) <https://www.scmp.com/news/china/society/article/3035983/chinese-primary-school-stops-using-headbands-study-pupils> accessed 17 September 2026.
[307] Kiesow Cortez (n 35).

Institutional design constitutes the final component of the framework. The European Union already possesses supervisory authorities under the General Data Protection Regulation, market-surveillance structures under the Artificial Intelligence Act and the Medical Devices Regulation, and new arrangements for health-data governance under the European Health Data Space. The United States distributes authority across the Federal Trade Commission, the Department of Health and Human Services, the Food and Drug Administration, state attorneys general, state privacy agencies, and a range of sectoral institutions. A framework governing neural data does not require the creation of a single new super-regulator, but it does require coordination and technical capacity of a kind not presently in evidence. Regulators must be in a position to inspect data-flow manifests, to test claims of local processing, to reproduce differential-privacy accounting, to evaluate secure-aggregation guarantees, and to determine whether a system update has enabled a new category of inference.

A 2026 event supported by the European Brain Council proposed the funding of ethical innovation, regulatory cohesion, and public engagement,[308] including recommendations on the establishment of a Neurotechnology Funding Board to navigate and adapt existing European funding instruments, to support innovation hubs, and to pool reimbursement schemes across Member States so that companies may scale across the single market. It was also recommended to employ the Digital Omnibus so as to recognize brain data as sensitive data under the General Data Protection Regulation, together with cohesion as between the medical and consumer device frameworks, regulatory sandboxes, and improved pathways for early dialogue with regulators. Further recommendations included a call for proactive lifecycle engagement with users and with

[308] European Brain Council (EBC) (n 115).

civil society, investment in interdisciplinary skills and common terminology, and a central role for soft-law instruments including the Charter[309] adopted by the European Brain Council.

The debate concerning the proper object of protection is of particular relevance to the argument advanced here. If emotion, attention, or intention may be inferred from both neural and nonneural signals, then regulating only the source of the signal risks producing arbitrary gaps in protection.[310] At the EBC event it was argued that that the gap within the General Data Protection Regulation concerns inferences about the mind rather than neurotechnology as such, and asking why an emotion should attract higher protection by reason of having been processed through a neurodevice when the same emotion inferred from a text-based technology attracts none.[311] It was also observed that the Artificial Intelligence Act may matter more than the General Data Protection Regulation for the governance of neurotechnology, bearing directly upon the analysis in Part IV above. Representation from the European Commission also argued that a dedicated Neurotech Act is not viewed as a requirement, the task being instead to build upon existing foundations.[312] Soft law was viewed as not a substitute for enforceable rights, but it was acknowledged that it would be capable of supplying updateable technical and ethical benchmarks more rapidly than legislation acting alone.[313]

Transatlantic harmonization ought accordingly to concentrate upon interoperable evidence rather than upon identical statutes.[314] The European Union and the United States are unlikely to converge upon a single privacy code within the near term, and it would be unrealistic to propose that they

[309] European Brain Council (EBC) (n 136).
[310] Magee, Ienca and Farahany (n 88).
[311] European Brain Council (EBC) (n 115).
[312] ibid.
[313] ibid.
[314] Kiesow Cortez (n 35).

should. They are nevertheless capable of recognizing common technical evidence in respect of confidentiality, minimization, controlled computation, secure deletion, and auditability.[315] A manufacturer selling a neural device upon both markets ought to be in a position to demonstrate, by means of a single technically rigorous dossier, what data leaves the device, which party is able to decrypt it, what model-update leakage has been tested, what secondary inferences are technically blocked, and how the system is re-evaluated following change. Jurisdiction-specific law would continue to determine whether the underlying purpose is permissible.

International instruments already supply the shared principles upon which such convergence might be constructed, and the trajectory of their development is itself instructive, disclosing a steady movement from voluntary innovation guidance towards formal normative standard-setting. The Organisation for Economic Co-operation and Development began the process in 2019 with its Recommendation on Responsible Innovation in Neurotechnology, the first intergovernmental standard in the field, which framed its concerns in the language of responsible innovation rather than of rights.[316] The Council of Europe subsequently assessed the implications of neurotechnology for the existing human-rights and data-protection framework.[317] The movement culminated, at least for the present, in the Recommendation on the Ethics of Neurotechnology adopted by UNESCO in 2025[318], the first genuinely global standard-setting instrument in the field, which consolidated the principles dispersed across the earlier texts into a single framework addressed to states. The direction of travel is therefore from soft and sectoral guidance towards

[315] Near and others (n 275).
[316] OECD (n 7).
[317] Bertoni and Ienca (n 141).
[318] UNESCO (n 8).

comprehensive and universal instruments, and it is that convergence of principle which renders a coordinated regulatory response conceivable rather than merely aspirational.

Professional and sectoral bodies have begun to align themselves with the same principles. The American Medical Association has endorsed the adoption of laws and regulations safeguarding neural data[319], and industry-facing initiatives have moved in parallel: the European Brain Council, having adopted its European Charter for the Responsible Development of Neurotechnologies in April 2025, launched a Seal of Responsible Neurotechnology as a voluntary mechanism by which organizations might demonstrate, in a verifiable manner, their alignment with the principles of the Charter.[320] The significance of such a mechanism is that it operates upon the supply side, offering a means by which responsible actors may distinguish themselves within precisely the consumer market whose disclosure practices have been shown to be least adequate, and thereby supplying through reputational incentive what formal regulation of the wellness sector has so far failed to supply through obligation.

A standards-based compliance pathway is capable of reducing that uncertainty where it creates reusable evidence, whereas a purely vague duty of reasonable security compels each firm to invent its own compliance narrative. The regulatory objective is not aiming for maximal restriction but by permitting socially valuable clinical and research uses under architectures which reduce unnecessary exposure, while prohibiting coercive, manipulative, or discriminatory uses.

[319] American Medical Association, 'Safeguarding Neural Data Collected by Neurotechnologies H-315.957' (2025) Resolution 503 A-25 <https://policysearch.ama-assn.org/policyfinder/detail/H-315.957?uri=%2FAMADoc%2FHOD.xml-H-315.957.xml>.
[320] European Brain Council (EBC) (n 115).

XIII. Conclusion

Neurotechnology complicates the protection of mental privacy because it renders imprecise the boundary between physiological measurement, personal data, health information, and mental inference. The legal response cannot rest upon a single claim that neural data is uniquely sensitive. Source-based definitions are useful insofar as they bring obvious measurements of the brain and nervous system within heightened protection, but they remain vulnerable to the same boundary problems which attended genetic and biometric exceptionalism. The more robust approach regulates the combination of inference, purpose, sector, power asymmetry, and architecture.

The comparison undertaken in this paper discloses different institutional strengths within the two jurisdictions examined. European law supplies a horizontal rights-based framework through the General Data Protection Regulation, context-specific restrictions through the Artificial Intelligence Act, product oversight through the Medical Devices Regulation, and emerging architectural governance through the European Health Data Space. American law supplies strong but fragmented sectoral regimes, consumer-protection authority vested in the Federal Trade Commission, medical-device oversight exercised by the Food and Drug Administration, and an expanding patchwork of state legislation. Colorado, California, Montana, and Connecticut demonstrate a shared legislative intuition that neural data merits heightened treatment, while their divergent definitions demonstrate why interoperability at the national level remains difficult to achieve. The Vermont enactments of 2026 confirm that the patchwork continues to grow, even where, as in the case of Act 101, the operative content was legislated away in the course of passage.

The central recommendation advanced here is that legal standards be rendered technically reviewable. Encryption should constitute a baseline in respect of identifiable neural information, though not a complete safe harbour. Local processing should constitute the presumptive

architecture wherever the user-facing function may be delivered without export of raw signal, Differential privacy should be considered in respect of aggregate analytics and of appropriate model-training outputs, with documented units, parameters, and cumulative accounting. Federated learning should not to be treated as inherently private and should be combined with secure aggregation and, where justified, with differential privacy. High-sensitivity cross-institutional computation should consider secure multiparty computation or homomorphic encryption. Above all, regulators should require controllers to document threat models, residual risks, prohibited inferences, update triggers, and independent testing.

Mental privacy is therefore best protected neither by the language of rights alone nor by technology alone. Cognitive liberty, mental integrity, and freedom of thought are capable of providing normative principles; sector-specific law is capable of determining when consent is meaningful and when a use should be prohibited; institutional oversight is capable of supplying enforcement; and privacy-enhancing technologies are capable of converting abstract obligations into properties which engineers and auditors are able to test. The strongest form of privacy by design is achieved where systems are constructed not merely so as to protect a database following collection, but so as to minimize what becomes legible, transferable, and reusable in the first place.

This paper has sought to show that the debate over whether neural data constitute a special category has been conducted upon the wrong axis. The exceptionalism question, which has absorbed so much of the literature, turns out to matter far less than the questions which lie beneath it: what may be inferred from a given signal, in what context, and by whom; where the processing occurs; and whether the abstract legal standards of security, minimization, and protection by design can be rendered technically reviewable. The experience of genetic privacy suggests that a category, however carefully drawn, might do little work on its own, and that the difficulties which arise are

difficulties of coverage and of enforcement rather than of definition. The empirical survey of decoding capability suggests that regulation solely focusing on present capability might be quickly overtaken, while the architecture of the systems being built today will matter for better safeguards. And the two bodies of law examined here, the European and the American, though they proceed from different premises, converge upon the same unmet need: a means of translating what the law requires in the abstract into how a neurotechnology device must be build and how an impact assessment audit can be conducted.

One way forward follows from these findings is to attach onto the existing frameworks a graded and inferential-risk-based obligation, adjusted to what a system can infer rather than to where its signals originate; to make processing locality the default and transmission the exception requiring justification; to require that safeguards be documented, combined against an explicit threat model, and revised through an assessment which a supervisory authority can audit; and to attach these requirements to reusable models and infrastructures across their lifecycle rather than only to individual devices at the point of sale. Given the importance of the domain regarding fundamental rights, what will require the attention of regulators on both sides of the Atlantic and further debate is to specify instruments with enough precision that compliance becomes reviewable and evasion becomes more visible.

References

Acquisti A, Taylor C and Wagman L, ‘The Economics of Privacy’ (2016) 54 Journal of Economic Literature 442 <https://doi.org/10.1257/jel.54.2.442>

Agarwal A and others, ‘Protecting Privacy of Users in Brain-Computer Interface Applications’ (2019) 27 IEEE Transactions on Neural Systems and Rehabilitation Engineering 1546 <https://doi.org/10.1109/TNSRE.2019.2926965>

American Medical Association, ‘Safeguarding Neural Data Collected by Neurotechnologies H-315.957’ (2025) Resolution 503 A-25 <https://policysearch.ama-assn.org/policyfinder/detail/H-315.957?uri=%2FAMADoc%2FHOD.xml-H-315.957.xml>

‘Are Your Thoughts Your Own?:"Neuroprivacy" and the Legal Implications of Brain Imaging’ (*New York City Bar Association*) <https://www.nycbar.org/reports/are-your-thoughts-your-ownneuroprivacy-and-the-legal-implications-of-brain-imaging/> accessed 17 September 2026

Barrett LF and others, ‘Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements’ (2019) 20 Psychological Science in the Public Interest 1 <https://doi.org/10.1177/1529100619832930>

Bertoni E and Ienca M, ‘The Privacy and Data Protection Implication of the Use of Neurotechnology and Neural Data from the Perspective of Convention 108’ (Council of Europe 2024)

Birbaumer N, ‘Breaking the Silence: Brain–Computer Interfaces (BCI) for Communication and Motor Control’ (2006) 43 Psychophysiology 517 <https://doi.org/10.1111/j.1469-8986.2006.00456.x>

Bonaci T, Calo R and Chizeck HJ, ‘App Stores for the Brain : Privacy and Security in Brain-Computer Interfaces’ (2015) 34 IEEE Technology and Society Magazine 32 <https://doi.org/10.1109/MTS.2015.2425551>

Brauneck A and others, ‘Federated Machine Learning, Privacy-Enhancing Technologies, and Data Protection Laws in Medical Research: Scoping Review’ (2023) 25 Journal of Medical Internet Research e41588 <https://doi.org/10.2196/41588>

Bublitz JC, ‘Novel Neurorights: From Nonsense to Substance’ (2022) 15 Neuroethics 7 <https://doi.org/10.1007/s12152-022-09481-3>

Card NS and others, ‘Long-Term Independent Use of an Intracortical Brain–Computer Interface for Speech and Cursor Control’ (2026) 32 Nature Medicine 2504 <https://doi.org/10.1038/s41591-026-04414-6>

CFG, ‘Neurotech Consumer Market Atlas - How the Sector Is Making Moves into the Mainstream’ (*Centre for Future Generations*, 16 June 2025) <https://cfg.eu/neurotech-market-atlas/> accessed 16 September 2026

——, ‘Towards an EU Neurotechnology Strategy’ (*Centre for Future Generations*, 10 April 2026) <https://cfg.eu/towards-an-eu-neurotechnology-strategy/> accessed 16 September 2026

‘Chile: Pioneering the Protection of Neurorights | The UNESCO Courier’ <https://courier.unesco.org/en/articles/chile-pioneering-protection-neurorights> accessed 14 September 2026

‘Chinese School Forced to Stop Using Headbands to Monitor Concentration’ (*South China Morning Post*, 1 November 2019) <https://www.scmp.com/news/china/society/article/3035983/chinese-primary-school-stops-using-headbands-study-pupils> accessed 17 September 2026

Clark, Linda, Crespo, Melissa and Wang, Katherine, ‘Neural Data Added to Montana’s Genetic Information Privacy Act’ (*Morrison Foerster*, August 2025) <https://www.mofo.com/resources/insights/null> accessed 1 October 2025

Clayton EW and others, ‘The Law of Genetic Privacy: Applications, Implications, and Limitations’ (2019) 6 Journal of Law and the Biosciences 1 <https://doi.org/10.1093/jlb/lsz007>

Cohen A and Nissim K, ‘Towards Formalizing the GDPR’s Notion of Singling out’ (2020) 117 Proceedings of the National Academy of Sciences of the United States of America 8344

Computer Security Division ITL, ‘Privacy-Enhancing Cryptography | CSRC | CSRC’ (*CSRC | NIST*, 3 January 2017) <https://csrc.nist.gov/projects/pec> accessed 17 September 2026

Cornejo-Plaza MI, Cippitani R and Pasquino V, ‘Chilean Supreme Court Ruling on the Protection of Brain Activity: Neurorights, Personal Data Protection, and Neurodata’ (2024) 15 Frontiers in Psychology 1330439 <https://doi.org/10.3389/fpsyg.2024.1330439>

Défossez A and others, ‘Decoding Speech Perception from Non-Invasive Brain Recordings’ (2023) 5 Nature Machine Intelligence 1097 <https://doi.org/10.1038/s42256-023-00714-5>

Dwork C, ‘Differential Privacy: A Survey of Results’ in Manindra Agrawal and others (eds), *Theory and Applications of Models of Computation* (Springer 2008) <https://doi.org/10.1007/978-3-540-79228-4_1>

European Brain Council (EBC), ‘Event Highlights: A European Strategy for Neurotechnology’ (20 April 2026) <https://www.braincouncil.eu/event-highlights-a-european-strategy-for-neurotechnology/> accessed 16 September 2026

——, ‘European Charter for the Responsible Development of Neurotechnologies’ (23 September 2026) <https://www.braincouncil.eu/european-charter-for-the-responsible-development-of-neurotechnologies/> accessed 16 September 2026

European Data Protection Board, ‘Guidelines 4/2019 on Article 25 Data Protection by Design and by Default’ (2020)

——, ‘Guidelines 01/2025 on Pseudonymisation’ (2025)

——, ‘Guidelines 1/2026 on Processing of Personal Data for Scientific Research Purposes | European Data Protection Board’ (25 June 2026) <https://www.edpb.europa.eu/public-consultations/guidelines-12026-on-processing-of-personal-data-for-scientific-research_en> accessed 17 September 2026

European Group on Ethics in Science and New Technologies, ‘Governing Neuro-AI: Towards an Infrastructure Approach’ (2026)

European Parliament. Directorate General for Parliamentary Research Services., *The Protection of Mental Privacy in the Area of Neuroscience: Societal, Legal and Ethical Challenges.* (Publications Office 2024) <https://data.europa.eu/doi/10.2861/869928> accessed 30 September 2025

Farahany N, ‘The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology’ [2023] Faculty Books

Farahany NA, ‘Incriminating Thoughts’ (2012) 64 Stanford Law Review 351

‘FTC Enforcement Action to Bar GoodRx from Sharing Consumers’ Sensitive Health Info for Advertising’ (*Federal Trade Commission*, 31 January 2023) <https://www.ftc.gov/news-events/news/press-releases/2023/02/ftc-enforcement-action-bar-goodrx-sharing-consumers-sensitive-health-info-advertising> accessed 17 September 2026

Galli G and others, ‘Early EEG Responses to Pre-Electoral Survey Items Reflect Political Attitudes and Predict Voting Behavior’ (2021) 11 Scientific Reports 18692 <https://doi.org/10.1038/s41598-021-96193-y>

Garrison NA and others, ‘Genomic Contextualism: Shifting the Rhetoric of Genetic Exceptionalism’ (2019) 19 The American Journal of Bioethics 51 <https://doi.org/10.1080/15265161.2018.1544304>

Genser, Jared, Damianos, Stephen and Yuste, Rafael, ‘Safeguarding Brain Data: Assessing the Privacy Practices of Consumer Neurotechnology Companies’ (Neurorights Foundation 2024)

Green RC, Lautenbach D and McGuire AL, ‘GINA, Genetic Discrimination, and Genomic Medicine’ (2015) 372 The New England Journal of Medicine 397 <https://doi.org/10.1056/NEJMp1404776>

‘Health Breach Notification Rule’ (*Federal Trade Commission*, 19 September 2014) <https://www.ftc.gov/legal-library/browse/rules/health-breach-notification-rule> accessed 17 September 2026

Health C for D and R, ‘Implanted Brain-Computer Interface (BCI) Devices for Patients with Paralysis or Amputation - Non-Clinical Testing and Clinical Considerations’ (17 May 2021) <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/implanted-brain-computer-interface-bci-devices-patients-paralysis-or-amputation-non-clinical-testing> accessed 17 September 2026

Hertz N, ‘Neurorights – Do We Need New Human Rights? A Reconsideration of the Right to Freedom of Thought’ (2022) 16 Neuroethics 5 <https://doi.org/10.1007/s12152-022-09511-0>

Ienca M, ‘On Neurorights’ (2021) 15 Frontiers in Human Neuroscience <https://doi.org/10.3389/fnhum.2021.701258> accessed 15 September 2026

——, ‘Towards a Governance Framework for Brain Data’ (2022) 15 Neuroethics 20 <https://doi.org/10.1007/s12152-022-09498-8>

Ienca M and Andorno R, ‘Towards New Human Rights in the Age of Neuroscience and Neurotechnology’ (2017) 13 Life Sciences, Society and Policy 5 <https://doi.org/10.1186/s40504-017-0050-1>

Ienca M and Haselager P, ‘Hacking the Brain: Brain–Computer Interfacing Technology and the Ethics of Neurosecurity’ (2016) 18 Ethics and Information Technology 117 <https://doi.org/10.1007/s10676-016-9398-9>

Ienca M and Malgieri G, ‘Mental Data Protection and the GDPR’ (2022) 9 Journal of Law and the Biosciences lsac006 <https://doi.org/10.1093/jlb/lsac006>

Istace T, ‘Establishing Neurorights: New Rights versus Derived Rights’ (2025) 17 Journal of Human Rights Practice 121 <https://doi.org/10.1093/jhuman/huae042>

Ivanković V, ‘What (If Anything) Morally Separates Environmental from Neurochemical Behavioral Interventions?’ (2023) 17 Neuroethics 6 <https://doi.org/10.1007/s12152-023-09540-3>

Jiang W-B, Zhao L-M and Lu B-L, ‘Large Brain Model for Learning Generic Representations with Tremendous EEG Data in BCI’ (arXiv, 29 May 2024) <https://doi.org/10.48550/arXiv.2405.18765> accessed 17 September 2026

Joly Y and others, ‘Looking Beyond GINA: Policy Approaches to Address Genetic Discrimination’ (2020) 21 Annual Review of Genomics and Human Genetics 491 <https://doi.org/10.1146/annurev-genom-111119-011436>

——, ‘Establishing the International Genetic Discrimination Observatory’ (2020) 52 Nature Genetics 466 <https://doi.org/10.1038/s41588-020-0606-5>

Joly Y and Dalpe G, ‘Genetic Discrimination Still Casts a Large Shadow in 2022’ (2022) 30 European Journal of Human Genetics 1320 <https://doi.org/10.1038/s41431-022-01194-8>

Jones CI and Tonetti C, ‘Nonrivalry and the Economics of Data’ (2020) 110 American Economic Review 2819 <https://doi.org/10.1257/aer.20191330>

Jurídica EA, ‘Corte Suprema ordenó eliminar información recogida por dispositivo de monitoreo de actividad cerebral del recurrente’ (*DOE | Actualidad Jurídica*, 11 August 2023) <https://actualidadjuridica.doe.cl/corte-suprema-ordeno-eliminar-informacion-recogida-por-dispositivo-de-monitoreo-de-actividad-cerebral-del-recurrente/> accessed 14 September 2026

Kiesow Cortez E, ‘Neurotechnology and Privacy: A Comparative Analysis of US and EU Approaches’ (TTLF Working Papers, Stanford-Vienna Transatlantic Technology Law Forum 2025)

Kousi H, ‘EBC Introduces the Seal of Responsible Neurotechnology’ (*European Brain Council (EBC)*, 9 April 2026) <https://www.braincouncil.eu/ebc-seal-of-responsible-neurotechnologies/> accessed 16 September 2026

Kunz EM and others, ‘Inner Speech in Motor Cortex and Implications for Speech Neuroprostheses’ (2025) 188 Cell 4658 <https://doi.org/10.1016/j.cell.2025.06.015>

Ligthart S, ‘Freedom of Thought in Europe: Do Advances in “Brain-Reading” Technology Call for Revision?’ (2020) 7 Journal of Law and the Biosciences lsaa048 <https://doi.org/10.1093/jlb/lsaa048>

——, ‘Minding Rights: Mapping Ethical and Legal Foundations of “Neurorights”’ (2023) 32 Cambridge Quarterly of Healthcare Ethics 461 <https://doi.org/10.1017/S0963180123000245>

Magee P, Ienca M and Farahany N, ‘Beyond Neural Data: Cognitive Biometrics and Mental Privacy’ (2024) 112 Neuron 3017 <https://doi.org/10.1016/j.neuron.2024.09.004>

McMahan HB and others, ‘Communication-Efficient Learning of Deep Networks from Decentralized Data’ (2017) Proceedings of the 20 th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017. JMLR: W&CP volume 54 <https://research.google/pubs/communication-efficient-learning-of-deep-networks-from-decentralized-data/> accessed 17 September 2026

Meng L and others, ‘Protecting Multiple Types of Privacy Simultaneously in EEG-Based Brain-Computer Interfaces’ (arXiv, 29 November 2024) <https://doi.org/10.48550/arXiv.2411.19498> accessed 17 September 2026

‘Mobile Health App Interactive Tool’ (*Federal Trade Commission*, 7 December 2022) <https://www.ftc.gov/business-guidance/resources/mobile-health-apps-interactive-tool> accessed 17 September 2026

Muhl E, ‘The Challenge of Wearable Neurodevices for Workplace Monitoring: An EU Legal Perspective’ (2024) 6 Frontiers in Human Dynamics <https://doi.org/10.3389/fhumd.2024.1473893> accessed 16 September 2026

Murphy ERD and Rissman J, ‘Evidence of Memory from Brain Data’ (2020) 7 Journal of Law and the Biosciences lsaa078 <https://doi.org/10.1093/jlb/lsaa078>

——, ‘Evidence of Memory from Brain Data’ (2020) 7 Journal of Law and the Biosciences lsaa078 <https://doi.org/10.1093/jlb/lsaa078>

Murray TH, ‘Genetic Exceptionalism and ’Future Diaries’ : Is Genetic Information Different from Other Medical Information?’ [1997] Genetic Secrets : Protecting Privacy and Confidentiality in the Genetic Era 60

——, ‘Is Genetic Exceptionalism Past Its Sell-By Date? On Genomic Diaries, Context, and Content’ (2019) 19 The American Journal of Bioethics 13 <https://doi.org/10.1080/15265161.2018.1552038>

Nasr M, Shokri R and Houmansadr A, ‘Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-Box Inference Attacks against Centralized and Federated Learning’ *2019 IEEE Symposium on Security and Privacy (SP)* (2019) <https://doi.org/10.1109/SP.2019.00065> accessed 17 September 2026

Near J, ‘Protecting Model Updates in Privacy-Preserving Federated Learning’ [2024] NIST <https://www.nist.gov/blogs/cybersecurity-insights/protecting-model-updates-privacy-preserving-federated-learning> accessed 17 September 2026

——, ‘Guidelines for Evaluating Differential Privacy Guarantees’ (National Institute of Standards and Technology 2025) NIST Special Publication (SP) 800-226 <https://doi.org/10.6028/NIST.SP.800-226> accessed 16 September 2026

Nissim K and others, ‘Bridging the Gap between Computer Science and Legal Approaches to Privacy’ (2017) 31 Harvard Journal of Law & Technology (Harvard JOLT) 687

‘NIST Privacy Framework Version 1.0’ [2020] NIST <https://www.nist.gov/privacy-framework> accessed 16 September 2026

OECD, ‘Recommendation of the Council on Responsible Innovation in Neurotechnology’ (12 November 2019) <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0457> accessed 14 September 2026

Ohm P, ‘Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization’ (2010) 57 UCLA Law Review 1701

Patrick-Krueger KM, Burkhart I and Contreras-Vidal JL, ‘The State of Clinical Trials of Implantable Brain–Computer Interfaces’ (2025) 3 Nature Reviews Bioengineering 50 <https://doi.org/10.1038/s44222-024-00239-5>

——, ‘The State of Clinical Trials of Implantable Brain–Computer Interfaces’ (2025) 3 Nature Reviews Bioengineering 50 <https://doi.org/10.1038/s44222-024-00239-5>

Paul S and Bajwa G, ‘Q-DP-GAN: Improving EEG Data Privacy Through Quantum-Inspired Differential Privacy-Based GAN’ (2026) 10 Cryptography 31 <https://doi.org/10.3390/cryptography10030031>

Popescu AB and others, ‘Privacy Preserving Classification of EEG Data Using Machine Learning and Homomorphic Encryption’ (2021) 11 Applied Sciences 7360 <https://doi.org/10.3390/app11167360>

‘Protecting the Mental Realm: What Does Human Rights Law Bring to the Table? - Timo Istace, 2023’ <https://journals.sagepub.com/doi/10.1177/09240519231211823> accessed 14 September 2026

Quan, ‘Neural Data Regulation: Lessons from Biometric Data’s Evolution’ (*California Lawyers Association*, 15 May 2026) <https://calawyers.org/privacy-law/neural-data-regulation-lessons-from-biometric-datas-evolution/> accessed 17 September 2026

Rainey S and others, ‘Is the European Data Protection Regulation Sufficient to Deal with Emerging Data Concerns Relating to Neurotechnology?’ (2020) 7 Journal of Law and the Biosciences lsaa051 <https://doi.org/10.1093/jlb/lsaa051>

Rothstein MA, ‘Genetic Exceptionalism and Legislative Pragmatism’ (2007) 35 Journal of Law, Medicine & Ethics 59 <https://doi.org/10.1111/j.1748-720X.2007.00154.x>

Rubinstein IS and Good N, ‘The Trouble with Article 25 (and How to Fix It): The Future of Data Protection by Design and Default’ (2020) 10 International Data Privacy Law 37 <https://doi.org/10.1093/idpl/ipz019>

——, ‘The Trouble with Article 25 (and How to Fix It): The Future of Data Protection by Design and Default’ (2020) 10 International Data Privacy Law 37 <https://doi.org/10.1093/idpl/ipz019>

Rubinstein IS and Hartzog W, ‘Anonymization and Risk’ (2016) 91 Washington Law Review 703

Ruiz S and others, ‘Neurorights in the Constitution: From Neurotechnology to Ethics and Politics’ (2024) 379 Philosophical Transactions of the Royal Society B: Biological Sciences 20230098 <https://doi.org/10.1098/rstb.2023.0098>

——, ‘Neurorights in the Constitution: From Neurotechnology to Ethics and Politics’ (2024) 379 Philosophical Transactions of the Royal Society B: Biological Sciences 20230098 <https://doi.org/10.1098/rstb.2023.0098>

‘Sens. Cantwell, Schumer, Markey Introduce Legislation to Shield Americans’ Brain Data From Exploitation’ (*U.S. Senate Committee on Commerce, Science, & Transportation*) <https://www.commerce.senate.gov/press/dem/release/sens-cantwell-schumer-markey-introduce-legislation-to-shield-americans-brain-data-from-exploitation/> accessed 17 September 2026

Shirakawa K and others, ‘Spurious Reconstruction from Brain Activity’ (2025) 190 Neural Networks 107515 <https://doi.org/10.1016/j.neunet.2025.107515>

‘Supreme Court of Chile, Role 105.065-2023, Judgment of 9-8-2023’ <https://www.doe.cl/alerta/11082023/20230811001>

Susser D and Cabrera LY, ‘Brain Data in Context: Are New Rights the Way to Mental and Brain Privacy?’ (2024) 15 AJOB Neuroscience 122 <https://doi.org/10.1080/21507740.2023.2188275>

——, ‘Brain Data in Context: Are New Rights the Way to Mental and Brain Privacy?’ (2024) 15 AJOB Neuroscience 122 <https://doi.org/10.1080/21507740.2023.2188275>

Suter SM, *The Allure and Peril of Genetics Exceptionalism: Do We Need Special Genetics Legislation?* (2001)

Szoszkiewicz Ł and Yuste R, 'Mental Privacy: Navigating Risks, Rights and Regulation : Advances in Neuroscience Challenge Contemporary Legal Frameworks to Protect Mental Privacy' (2025) 26 EMBO reports 3469 <https://doi.org/10.1038/s44319-025-00505-6>

Tang J and others, 'Semantic Reconstruction of Continuous Language from Non-Invasive Brain Recordings' (2023) 26 Nature Neuroscience 858 <https://doi.org/10.1038/s41593-023-01304-9>

'TechDispatch #1/2024 - Neurodata | European Data Protection Supervisor' (1 October 2025) <https://www.edps.europa.eu/data-protection/our-work/publications/techdispatch/2024-06-03-techdispatch-12024-neurodata> accessed 2 October 2025

'——' (16 August 2026) <https://www.edps.europa.eu/data-protection/our-work/publications/techdispatch/2024-06-03-techdispatch-12024-neurodata> accessed 14 September 2026

Technology NI of S and, 'Advanced Encryption Standard (AES)' (US Department of Commerce 2023) Federal Information Processing Standard (FIPS) 197 <https://doi.org/10.6028/NIST.FIPS.197-upd1> accessed 17 September 2026

Tesink V and others, 'Right to Mental Integrity and Neurotechnologies: Implications of the Extended Mind Thesis' (2024) 50 Journal of Medical Ethics 656 <https://doi.org/10.1136/jme-2023-109645>

'The "Neural Data" Goldilocks Problem: Defining "Neural Data" in U.S. State Privacy Laws' (*https://fpf.org/*) <https://fpf.org/blog/the-neural-data-goldilocks-problem-defining-neural-data-in-u-s-state-privacy-laws/> accessed 15 September 2026

'The Protection of Mental Privacy in the Area of Neuroscience - Societal, Legal and Ethical Challenges | Think Tank | European Parliament' <https://www.europarl.europa.eu/thinktank/en/document/EPRS_STU(2024)757807> accessed 14 September 2026

UNESCO, 'Recommendation on the Ethics of Neurotechnology - Legal Affairs' (11 November 2025) <https://www.unesco.org/en/legal-affairs/recommendation-ethics-neurotechnology> accessed 14 September 2026

Waldman AE, 'Data Protection by Design? A Critique of Article 25 of the GDPR' [2021] Cornell International Law Journal, Vol. 53, 2021

Wexler A and Reiner PB, 'Oversight of Direct-to-Consumer Neurotechnologies' (2019) 363 Science 234 <https://doi.org/10.1126/science.aav0223>

——, 'Oversight of Direct-to-Consumer Neurotechnologies' (2019) 363 Science 234 <https://doi.org/10.1126/science.aav0223>

Willett FR and others, ‘A High-Performance Speech Neuroprosthesis’ (2023) 620 Nature 1031 <https://doi.org/10.1038/s41586-023-06377-x>

Wu S, Ramdas A and Wehbe L, ‘Brainprints: Identifying Individuals from Magnetoencephalograms’ (2022) 5 Communications Biology 852 <https://doi.org/10.1038/s42003-022-03727-9>

Xia K and others, ‘Privacy-Preserving Brain–Computer Interfaces: A Systematic Review’ (2023) 10 IEEE Transactions on Computational Social Systems 2312 <https://doi.org/10.1109/TCSS.2022.3184818>

Yan N and others, ‘Lightweight Privacy-Preserving Feature Extraction for EEG Signals Under Edge Computing’ (2024) 11 IEEE Internet of Things Journal 2520 <https://doi.org/10.1109/JIOT.2023.3292232>

Yang Y-Y and others, ‘Person-Identifying Brainprints Are Stably Embedded in EEG Mindprints’ (2022) 12 Scientific Reports 17031 <https://doi.org/10.1038/s41598-022-21384-0>

Yuste R, ‘Advocating for Neurodata Privacy and Neurotechnology Regulation’ (2023) 18 Nature Protocols 2869 <https://doi.org/10.1038/s41596-023-00873-0>

——, ‘Four Ethical Priorities for Neurotechnologies and AI’ (2017) 551 Nature 159 <https://doi.org/10.1038/551159a>