

The Mind Is Not a Data Point: Neural Privacy and the Failure of Cognitive Liberty

Ismum Al-Araf
ismumalaraf@gmail.com

ABSTRACT

Neural data, defined as the information generated by consumer neurotechnology devices such as electroencephalogram (EEG) headsets and brain-computer interfaces, differs categorically from other forms of biometric data. Unlike a fingerprint or a facial scan, a neural recording encodes pre-conscious cognitive content that the user never chose to express. Current privacy frameworks fail to protect that content for a shared structural reason. The Health Insurance Portability and Accountability Act (HIPAA), the European Union's General Data Protection Regulation (GDPR), and California's Senate Bill 1223 were each designed to regulate behavioral outputs rather than the substrate of cognition itself. A 2024 audit found that 29 of 30 consumer neurotechnology companies retained unrestricted access to user brain data, and no United States federal statute specifically governs the collection or use of that data.

This analysis draws on recent neuroscience research, legal commentary on the emerging patchwork of state-level regulation, and the scholarly literature on cognitive liberty, particularly the framework developed by Nita Farahany and the UNESCO Recommendation on the Ethics of Neurotechnology adopted in November 2025. Evidence was collected through a structured search of peer-reviewed publications, legal advisories, and institutional reports issued between 2023 and 2026. The central finding is that cognitive liberty is well established in normative terms but remains unprotected in operational ones. The strongest objection to that finding, Chile's 2023 Supreme Court ruling enforcing neural privacy under existing constitutional law, receives direct treatment in the counterarguments section. That ruling demonstrates the need for systematic legislation rather than the sufficiency of current frameworks. Closing the regulatory gap requires extending data protection to inferences derived from neural signals, mandating opt-in consent, and imposing deletion timelines calibrated to anticipated advances in artificial intelligence decoding capability.

INTRODUCTION

In 2023, researchers at the University of Texas reconstructed the semantic content of imagined speech directly from non-invasive brain scans (Tang et al, 2023). The reconstructed material consisted not of words spoken aloud but of thoughts, recovered from neural activity alone. No law in the United States

specifically covers what those researchers decoded. Consumer EEG headsets that monitor brain activity for wellness purposes can be purchased today, and 29 of the 30 companies that sell such devices retain unrestricted access to the neural data collected (NeuroRights Foundation, 2024). Most of those companies also reserve the right to transfer that data to third parties.

The conditions described above are not a near-future concern. A \$15 billion industry currently operates in a near-total vacuum. Neural data differs from other biometric data not merely in sensitivity but in kind, because neural recordings encode cognitive content that users never chose to express. Existing privacy frameworks, built to regulate behavioral outputs, structurally cannot reach the substrate of thought. The categorical distinctiveness of neural data examines the failure of three major legal frameworks, assesses the gap between normative rights and enforcement mechanisms, and addresses the principal objections to that account.

METHODS

Evidence supporting the analysis was gathered through a structured literature search rather than through original empirical work. Three categories of source were targeted. Peer-reviewed neuroscience and neuroethics publications were identified through PubMed Central and journal databases using the search terms “neural data privacy,” “mental privacy,” “neurorights,” and “brain-computer interface decoding,” restricted to work published from 2023 onward. Legal and regulatory sources were identified through published advisories from law firms specializing in data privacy, as well as the statutory texts of HIPAA, the GDPR, and California Senate Bill 1223. Industry practice data was drawn from the 2024 audit of 30 consumer neurotechnology companies conducted by the NeuroRights Foundation, which remains the only systematic public assessment of privacy policies across the consumer sector.

Sources were selected against three criteria. Recency was prioritized because both decoding capability and regulatory activity changed substantially after 2023. Primary material was preferred over secondary summary, so original studies and statutory texts were consulted directly wherever available. Relevance to consumer rather than clinical neurotechnology governed inclusion, since clinical applications operate under separate regulatory regimes and raise different questions of consent. One study is cited at second hand, and that citation is marked as such in the text. The resulting analysis is a synthesis of published research, statutory language, and legal commentary current as of May 2026.

NEURAL DATA AS A DISTINCT CATEGORY OF INFORMATION

Biometric data describes the surface of the body. Fingerprints, retinal scans, and facial geometry record physical features that identify a person, capturing what a body is rather than what a mind contains. Neural data operates differently. When researchers at the University of Texas combined functional magnetic

resonance imaging (fMRI) recordings with a large language model, the resulting decoder recovered the gist of imagined narratives that participants had never spoken aloud (Tang et al., 2023). The recovered material was not identification of a person. The recovered material was cognition itself, rendered legible.

The distinction between gist-level and verbatim decoding determines how much the finding actually establishes. The decoder did not produce a word-for-word transcript of a participant's thoughts. The system instead generated continuous language that preserved the meaning of what a participant had heard, watched, or silently imagined, frequently using different vocabulary from the original. A stimulus describing a person rising and looking through a bedroom window, for example, was decoded as a description of walking toward a window and opening the glass: the event survives while the wording changes (Tang et al., 2023). Tang and colleagues reported performance above chance for perceived speech, for imagined speech, and for silent film, which indicates that the recovered representation was conceptual rather than acoustic or lexical.

Gist-level decoding is in one respect less alarming than verbatim decoding, since exact sentences remain private. In another respect, gist-level decoding is more alarming, because meaning rather than phrasing is what an advertiser, an employer, or a government would seek to extract. A system that reliably identifies what a person is thinking about, without reproducing the precise words, is sufficient for profiling, targeting, and prediction. The commercial value of neural data lies in semantic content, not in transcription.

Implanted brain-computer interfaces have advanced further than non-invasive methods. In 2024, an implanted device decoded the attempted speech of patients with amyotrophic lateral sclerosis at 97.5% accuracy (Card et al., 2024, as cited in Szoszkiewicz & Yuste, 2025). That technology requires neurosurgery and remains confined to clinical settings. Consumer EEG headsets, marketed as focus aids and sleep trackers, require no surgery at all. Scalp-level recordings from such devices already support inferences about inner language, attention, emotional state, and physiological arousal (Szoszkiewicz & Yuste, 2025)

Five properties distinguish neural data from every other category of personal information. Neural signals are captured involuntarily, because a person cannot choose which brain activity a sensor records. Neural signals reflect states that form before conscious decision, which places those signals outside the model of voluntary disclosure that privacy law assumes. Neural recording streams continuously rather than producing discrete events, so a single act of consent extends indefinitely. The consequences of extraction are irreversible, because cognitive content cannot be un-decoded once recovered. Neural data is finally constitutive of identity in a way that no purchase history or GPS trace achieves.

When a device registers a person's emotional state before that person has acted on the emotion, no act of sharing has taken place. The state was taken rather than given. Privacy law offers no framework for that transaction. Privacy law was not built for data that precedes the decision to disclose.

THE LEGAL FRAMEWORKS THAT SHOULD PROTECT NEURAL DATA CANNOT REACH IT

The scale of the regulatory gap becomes clear when measured against current industry practice. A 2024 audit of 30 consumer neurotechnology companies conducted by the NeuroRights Foundation found that 29 retained unrestricted access to user neural data. Among the companies surveyed, 96.7% reserved the right to transfer that data to third parties, fewer than 20% mentioned encryption, and only 16.7% committed to notifying users in the event of a breach (NeuroRights Foundation, 2024). No existing legal framework constrains any of those practices.

HIPAA protects neural data held by healthcare providers and health plans. Consumer wellness companies are not covered entities under the statute and therefore fall entirely outside its scope (Iliopoulos & Perkins, 2025). The GDPR classifies neural data as sensitive only where the data is used to uniquely identify a person. That threshold excludes the full range of cognitive-content applications, including emotional profiling, attention monitoring, neuromarketing, and behavioral prediction. A company that infers a user's emotional response to an advertisement is not identifying that user and may therefore operate outside the heightened protections of the GDPR.

California Senate Bill 1223, signed into law in September 2024, is the most protective neural data statute in the United States. The statute defines protected neural data as information generated by measuring nervous system activity that is not inferred from non-neural information. Legal analysts at Morrison Foerster have observed that the definition creates a structural gap (Clark & Martinez, 2024). If an EEG headset captures an emotional state and that inference is subsequently processed alongside other information, the resulting secondary data may fall outside the scope of the statute.

The most commercially valuable applications of neural data operate precisely at that secondary layer. Inferences that can be monetized, sold to advertisers, or used to profile psychological vulnerability sit one level above the raw signal, where no current law reaches. Federal action has not closed the gap. In September 2025, the strongest federal response was the Management of Individuals' Neural Data Act, which directs the Federal Trade Commission to study the issue. A directive to study a problem is not a protection against it.

COGNITIVE LIBERTY EXISTS IN PRINCIPLE BUT LACKS AN ENFORCEMENT MECHANISM

Scholarly work on neurotechnology has converged on a shared normative framework. Nita Farahany defines cognitive liberty as the master right of the neurotechnology era, comprising both a negative right that protects individuals from unwanted neural intrusion and a positive right that secures the freedom to understand one's own cognitive processes (Farahany, 2023). Across the three major neurorights

frameworks proposed in the academic literature, mental privacy is the principal point of consensus (Brown, 2024). International bodies have adopted comparable language. In November 2025, the 194 member states of UNESCO adopted a Recommendation on the Ethics of Neurotechnology, declaring the inviolability of the human mind, a foundation principle of global governance (O'Hagan & Parlier, 2025; Smedley, 2025).

The UNESCO agreement is non-binding. Article 18 of the International Covenant on Civil and Political Rights has protected freedom of thought since 1966, and that protection has not prevented the industry practices documented. The problem is not the absence of a principle but the absence of a mechanism to enforce it. Freedom of thought was conceived to protect the expression of belief, securing the right to dissent, to worship, and to speak without persecution. The doctrine was never designed to prevent a device from reading the neural substrate of cognition before a belief had been formed or expressed.

Cognitive liberty consequently exists in philosophical and normative terms while lacking the enforcement mechanisms that would make the right operational. Five such mechanisms are absent. Consumer neurotechnology companies face no data minimization requirements. No mandatory opt-in consent standard governs neural data collection. Legal protection extends to raw signals but not to cognitive inferences, and no deletion timelines prevent retroactive exploitation as decoder technology improves. Meaningful penalties for commercial violation do not exist. A right unaccompanied by such mechanisms functions as a statement of intention rather than as a protection.

COUNTERARGUMENTS

Three substantial objections can be raised against the account developed above. The first holds that existing constitutional protections are adequate where courts are willing to apply them. The second holds that current decoding capability has been overstated, which makes sector-specific regulation premature. The third holds that narrowly targeted regulation would impede medically beneficial innovation. Each objection is treated in turn below.

Existing Constitutional Protections are Adequate

The most serious objection originates in Chile. In 2023, the Chilean Supreme Court issued a unanimous ruling ordering Emotiv, a United States neurotechnology company, to erase brain data collected from former Senator Guido Girardi. The ruling rested on the existing constitutional right to privacy, supplemented by a 2021 constitutional amendment protecting neural data specifically (Spechak, 2025). No new neurorights statute was required to reach that outcome. If courts can enforce cognitive privacy through existing frameworks once cases reach them, the deficiency may lie in enforcement rather than in the substance of the law.

That objection carries weight, but three limitations constrain how far the Chilean example generalizes. First, the ruling depended on conditions absent elsewhere: a constitution amendment addressing neural data directly, the resources required to litigate to the nation's highest court, and the institutional standing of a former legislator. Most individuals whose brain data is collected by consumer devices possess none of those advantages. A legal protection available only to the privileged is not a protection at scale.

Second, the ruling arrived after the data had already been collected and processed. The precautionary principle, accepted even by scholars skeptical of new neurorights, requires that protection precede collection where the potential harm is irreversible. Neural data stored today under inadequate safeguards may be reprocessed by more capable decoders within a few years, meaning stored data has a future that current law has not anticipated. Third, Chile remains an exception rather than a template already in force. The United States has no equivalent constitutional provision, and most jurisdictions worldwide have nothing comparable. The Chilean case demonstrates that a constitutional commitment to neural privacy can be enforced, and replicating that commitment elsewhere requires deliberate legislation rather than judicial improvisation.

Current Decoding Capability Has Been Overstated

A second objection holds that regulatory urgency rests on an exaggerated account of what the technology can presently accomplish. The decoding results reported by Tang and colleagues required participants to spend many hours inside an fMRI scanner training a subject-specific model, and the resulting decoder performed poorly when applied to a different individual. Participants were also able to defeat the decoder deliberately by performing an unrelated mental task, which suggests that mental privacy retains a practical defense (Tang et al., 2023). Consumer EEG devices produce far noisier signals than fMRI and cannot approach comparable resolution.

Those technical constraints are real, but the constraints describe present capability rather than permanent limits. Decoder performance has improved as language models have grown more capable, and the cooperation requirement reflects current research protocols rather than a fixed property of the technology. The objection also addresses the wrong risk. Commercial harm does not require verbatim thought reading, because inferences about attention, emotional valence, and arousal are already commercially actionable and are already generated by consumer devices (Szozkiewicz & Yuste, 2025). Regulation calibrated to present capability alone would arrive too late, since data collected under today's standards will remain available to tomorrow's decoders.

Targeted Regulation Would Impede Beneficial Innovation

A third objection concerns cost. Neurotechnology offers substantial medical benefits, including restored communication for patients with paralysis and improved treatment options for epilepsy and severe depression. Sector-specific regulation may raise compliance costs, slow clinical translation, and push

development toward jurisdictions with weaker oversight. On this view, broad and technology-neutral privacy legislation would protect users without signaling out a nascent for special burden.

The concern is legitimate, and the proposals advanced here are drawn narrowly in response. Clinical neurotechnology already operates under HIPAA, institutional review board oversight, and Food and Drug Administration device regulation, so the additional requirements proposed above would fall primarily on consumer wellness products that currently face no comparable obligations. Technology-neutral privacy law has moreover had three decades in which to reach neural data and has not done so. The reason is structural: general frameworks defined around identification and behavioral output do not capture pre-behavioral cognitive content. A framework that has repeatedly failed to reach a category of data is unlikely to reach that category through further generalization.

LIMITATIONS

Several constraints limit the scope of the analysis presented here. The most significant constraint is the opacity of the industry under examination. Consumer neurotechnology companies are not required to disclose what neural data they collect, how long that data is retained, what inferences are derived from it, or to whom the data is transferred. The NeuroRights Foundation audit assessed published privacy policies rather than observed practice, which means the audit documents what companies reserve the right to do rather than what companies actually do. Actual practice may be either more protective or less protective than stated policy, and no public mechanisms exist to determine which.

A second constraint concerns proprietary capability. Decoding accuracy achieved inside private laboratories is not published, so the analysis relies on peer-reviewed results that may lag commercial capability by an unknown interval. Claims made about what can currently be inferred from consumer-grade EEG are therefore conservative by construction, and the true gap between capability and regulation may be wider than the published literature can demonstrate.

A third constraint is jurisdictional and temporal. The legal analysis concentrates on the United States, the European Union, and Chile, and does not assess frameworks in China, India, Brazil, or other jurisdictions with significant neurotechnology sectors. Regulatory activity in this area is also moving quickly, so the account of statutory law reflects conditions as of May 2026 and may be overtaken by subsequent legislation. No primary empirical research was conducted for analysis, which remains dependent on the accuracy and completeness of the published sources on which it draws.

CONCLUSION

In November 2025, 194 governments acknowledged that the relationship between technology and the self had fundamentally changed. That acknowledgement, recorded in a non-binding recommendation, has so

far produced no enforceable protection in most jurisdictions. The distance between what AI-powered neurotechnology can read from a brain scan and what existing privacy law covers is not a temporary lag. The gap is structural, because any framework that defines data protection in terms of behavioral output will continue to miss cognitive content, which by its nature precedes behavior.

Closing that gap requires four operational commitments: extending data protection to inferences derived from neural signals, requiring opt-in consent for neural data collection, imposing deletion timelines that account for advances in decoding, and bringing consumer neurotechnology companies within existing data protection regimes. None of those proposals is speculative. Each responds to an industry that is already operating, already collecting, and largely unaccountable. The contest over cognitive liberty is not approaching. The contest is already underway, and is already being lost quietly, in terms and conditions that few people read, through a technology that most people have not heard of, in a legal silence that policymakers have not yet chosen to fill.

REFERENCES

- Brown, Cohen, Marcus Lionel. "Neurorights, Mental Privacy, and Mind Reading." *Neuroethics*, vol. 17, no. 2, July 2024, <https://doi.org/10.1007/s12152-024-09568-z>.
- Clark, Linda K., and Carson Martinez. "A MoFo Privacy Minute Q&A: California Revises CCPA to Cover Neural Data." *Morrison Foerster*, 11 Oct. 2024, www.mofo.com/resources/insights/241011-a-mofo-privacy-minute-q-a-california-revises-ccpa.
- Farahany, Nita A. *The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology*. St. Martin's Press, 2023.
- Iliopoulos, Kristina, and Nancy L. Perkins. "Neural Data Privacy Regulation: What Laws Exist and What Is Anticipated?" *Arnold & Porter*, 2025, www.arnoldporter.com/en/perspectives/advisories/2025/07/neural-data-privacy-regulation.
- NeuroRights Foundation. *Safeguarding Brain Data: Assessing the Privacy Practices of Consumer Neurotechnology Companies*. NeuroRights Foundation, 2024.
- O'Hagan, Clare, and Romain Parlier. "Ethics of Neurotechnology: UNESCO Adopts the First Global Standard." *UNESCO*, 2025, www.unesco.org/en/articles/ethics-neurotechnology-unesco-adopts-first-global-standard-cutting-edge-technology.
- Smedley, Alison. "UNESCO Adopts First Global Framework on Neurotechnology Ethics." *UNESCO in the UK*, 18 Nov. 2025, unesco.org.uk/news/unesco-adopts-first-global-framework-on-neurotechnology-ethics.
- Spichak, Simon. "The Controversial Push for New Brain and Neurorights." *Journal of Medical Internet Research*, vol. 27, 25 Feb. 2025, e72270, <https://doi.org/10.2196/72270>.

Szoszkiewicz, Łukasz, and Rafael Yuste. "Mental Privacy: Navigating Risks, Rights and Regulation." *EMBO Reports*, vol. 26, no. 14, 25 June 2025, <https://doi.org/10.1038/s44319-025-00505-6>.

Tang, Jerry, et al. "Semantic Reconstruction of Continuous Language from Non-Invasive Brain Recordings." *Nature Neuroscience*, vol. 26, no. 5, 1 May 2023, pp. 858-866, <https://doi.org/10.1038/s41593-023-01304-9>.