NEURO-RIGHTS · POLICY & THE MIND

The Last Private Place

A Citizen's and Legislator's Guide to Neuro-Rights

For the first time in human history, the contents of a mind can be measured, decoded, nudged and sold. This is a guide — for citizens and for the legislators who will decide — to the rights that must guard the one frontier we have always assumed was our own.

A long read · policy & ethics · written for general readers and lawmakers alike

CONTENTS

1 The Frontier — Why the mind became a regulatory problem
2 Definitions — What we mean by "neuro-rights"
3 The Landscape — The law as it stands in 2026
4 Sovereignty of Thought — Guaranteeing the inner forum
5 Ownership — Can a company own your thoughts?
6 The Right Level — At what level should neuro-rights be legislated?
7 The Gaps — What key legislation is missing?
8 The Security Frontier — Neuroweapons & cognitive warfare
9 The Evidentiary Problem — Could you prove a mind was controlled?
10 The Benevolent Case — Good beyond helping the disabled
11 The Multiplier — How AI changes everything about neurotech
12 The Trade — What rights will we hand over?
13 Society — How everyday life changes
14 Victims — Defining, recognising and helping them
15 The Horizon — 3, 5, 10 and 20 years
16 Follow the Money — Who benefits commercially
17 The Classified Question — What might exist behind secrecy
18 Action — Immediate steps, by institution
19 Oversight — What it should look like
- In Closing — The window, and who holds it open
- References

Published as part of the TI One Voice initiative. TI One Voice (one.witysk.org) is an initiative founded by Stephane van der Aa to give a unified voice to the community of so-called "Targeted Individuals" (TIs) — people who report being subjected to organised harassment, surveillance, and what they describe as remote forms of physical and psychological abuse. The platform brings together their testimonies and supporting documentation in a single, structured place, so that experiences which are too often dismissed or scattered can be recorded, preserved, and examined seriously. Its goal is to document these accounts rigorously, to support the individuals affected, and to advocate before Belgian, European, and international institutions for recognition, protection, and proper investigation.

Every right we have ever written down assumed one thing: that no matter what was done to a person's body, their possessions, or their words, the inside of their head remained theirs. The mind was the one place a king could not enter, a court could not subpoena, and a thief could not reach. That assumption is now obsolete. Devices already on sale can read the electrical signature of attention, mood and intention. Laboratories can reconstruct words a person only imagined speaking, and images they only saw in their mind's eye. The question is no longer whether the mind can be touched from the outside. It is who will be allowed to touch it, under what conditions, and what happens to those who are touched without consent.

This essay is an attempt to map that question end to end — the law that exists, the law that is missing, the technology that is real, the technology that is hype, and the choices that fall to legislators in the next handful of years. It is deliberately long, because the subject does not reward shortcuts. Wherever the evidence is contested, I have said so plainly; on a topic this prone to both complacency and panic, honesty about what we do *not* know is part of the argument.

1. The Frontier — Why the mind became a regulatory problem

Neurotechnology is any tool that measures, interprets, modulates or stimulates the nervous system. That spans a vast range: the cochlear implant that has restored hearing to a million people; the deep-brain stimulator that quiets the tremor of Parkinson's disease; the EEG headband sold to improve sleep or focus; the experimental cortical implant that lets a paralysed person move a cursor by intention alone. What unites them is that they convert the activity of neurons into data, and increasingly run that data back into the brain.

Two forces turned this from a clinical curiosity into a governance emergency. The first is miniaturisation and commercialisation: brain-sensing has escaped the hospital and entered the consumer market in earbuds, headsets and wellness wearables. The second is artificial intelligence. Raw neural signals are noisy and, on their own, reveal little. Modern machine learning is extraordinarily good at finding the pattern in the noise — turning a smear of electrical activity into a reliable inference about what a person is feeling, attending to, or about to decide. Neurotechnology supplies the signal; AI supplies the meaning. Apart, each is limited. Together, they make the interior of a human being legible in a way no prior technology has approached.

In November 2025, the world's governments responded for the first time. Meeting in Samarkand, the General Conference of UNESCO adopted the **Recommendation on the Ethics of Neurotechnology** — the first global normative framework for the field — and framed its purpose in language that would have read as mysticism a generation ago: protecting the inviolability of the human mind.[1][2] That is the territory this essay covers.

> *"The mind is the seat of our identity, our memories, our intentions. If it can be read and rewritten by a third party, then every other right we hold is built on sand."*

— the premise of the neuro-rights movement

2. Definitions — What we actually mean by "neuro-rights"

The term *neurorights* was popularised by the neuroscientist Rafael Yuste of Columbia University and the interdisciplinary "Morningside Group," whose 2017 commentary in *Nature* argued that existing human-rights instruments were written in language too general to constrain brain technology, and proposed a small set of specific protections.[3][4] That framework now anchors most national legislation. It identifies five rights, shown below.

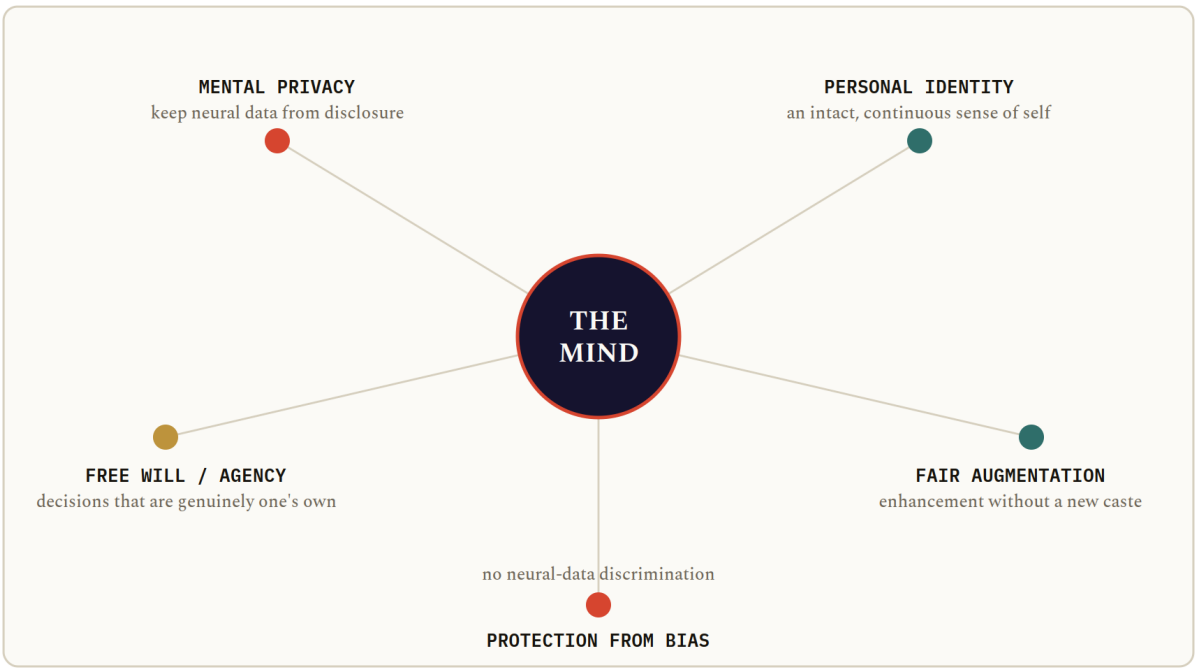


Figure 1. The five neurorights proposed by the Morningside Group (Yuste et al., Nature, 2017). An earlier, overlapping framework by Ienca and Andorno proposed four: cognitive liberty, mental privacy, mental integrity and psychological continuity.

It is worth being precise about a debate inside the field itself, because legislators will hear both sides. One camp argues these are genuinely *new* rights that must be written into law. A second camp — including some of the movement's own allies — argues that nothing new is needed, that mental privacy and integrity are already implied by existing rights to privacy, bodily integrity, freedom of thought and non-discrimination, and that the task is merely to *interpret* those rights for a new technology.[5] This is not a sterile academic quarrel. It determines whether a country needs a constitutional amendment, a new statute, or simply a regulator willing to apply the laws it already has.

3. The Landscape — The law as it stands in 2026

For a field this young, the legal landscape is moving fast. A useful way to read it is as a stack: a global recommendation at the top, a handful of pioneering national constitutions, a thickening layer of sub-national statutes, and a great deal of pending legislation underneath.



Figure 2. Eight years took the field from a journal commentary to the first global standard.

The UNESCO Recommendation (2025): the global keystone

Adopted on 12 November 2025 after six years of negotiation among UNESCO's member states, the Recommendation is the first attempt to set a rights-based standard covering the entire life cycle of neurotechnology — from how a product is designed, through deployment, to disposal.[1][2][6] It establishes high-level values (human dignity, mental privacy, freedom of thought, equity) and asks states to do concrete things: protect neural data, demand genuine and informed consent, guarantee transparency, and shield children and other vulnerable groups. Notably, it reaches beyond the clinic to address consumer and workplace uses, and on the contentious matter of *enhancement* it asks states to forbid any direct or indirect pressure to undergo neural augmentation, to prohibit enhancement that undermines dignity or equal opportunity, and to draw clear lines between what is acceptable, off-limits, or permitted only under strict oversight.[6]

> **The crucial limitation.** The Recommendation is **not legally binding**. It is, in the technical sense, "soft law": a standard states are urged to adopt, not a treaty they must follow. Its drafters note that UNESCO's 2021 AI-ethics recommendation helped shape the binding EU AI Act, so soft law can harden over time. But for now its force depends entirely on whether national legislatures act on it. That is precisely where the readers of this essay come in.

Chile: the country that wrote the mind into its constitution

Chile is the field's pioneer. On 14 October 2021, its Congress passed Law 21.383, amending Article 19 of the Constitution to require that scientific and technological development serve people and respect their physical and psychic integrity — the first time any nation gave brain activity explicit constitutional protection, with the word "neurorights" entering its legal vocabulary largely through Senator Guido Girardi and Rafael Yuste.[7][8]

The provision was tested almost immediately. In a ruling of 9 August 2023, Chile's Supreme Court found that the US consumer-neurotech company Emotiv had violated the constitutional rights to psychic integrity and privacy of Girardi himself, who had used the company's "Insight" EEG headset. The court ordered the firm to delete all of his brain data and rejected its argument that anonymising the data turned it into mere statistics exempt from consent.[8][9][10]

It was the world's first court decision squarely on neuro-privacy, and it proved that constitutional neurorights language could deliver a real remedy, not just a symbolic one.

> **What the Chilean case did not settle.** Legal scholars have since noted that the Emotiv ruling may have rested as much on ordinary data-protection law as on the new neurorights clause, and that *inferred* mental data — the conclusions an algorithm *draws* from your brain signals, as opposed to the raw signals themselves — still sits in a gray zone even under Chile's strong text.[11] That gap matters, because the inference is usually the part that can hurt you.

Latin America and Europe

Chile is not alone. Spain's 2021 Charter of Digital Rights includes neurorights provisions; Brazil moved to constitutionalise mental integrity; and Mexico, Colombia, Uruguay and Argentina have all seen legislative activity.[4][12] Europe, by contrast, has no dedicated neurorights instrument. It relies on a patchwork: the GDPR, which treats health and biometric data as special categories; the EU AI Act, which bans a narrow set of practices such as certain workplace and school emotion-recognition uses and some manipulative subliminal techniques; and the 2019 OECD Recommendation on Responsible Innovation in Neurotechnology, the first international standard of any kind here. The honest summary for a European legislator is that no current European law was written with brain-reading consumer devices in mind, and the fit is awkward.

The United States: a state-by-state mosaic

With no comprehensive federal privacy law, the United States has produced a patchwork at the state level. Colorado in 2024 became the first state to classify neural data as sensitive data under its privacy act; California amended its own law the same year; Montana and Connecticut followed in 2025; and at least fifteen further bills are pending.[13][14][15] At the federal level, the proposed MIND Act of 2025 would direct the Federal Trade Commission to study neural-data practices; the American Medical Association called for tighter regulation in mid-2025; and senators have pressed the FTC to investigate whether companies are exploiting consumers' brain data.[14][16]

> **The definition problem — a warning for drafters.** The US states cannot even agree on what "neural data" means. California excludes data *derived* from neural signals by an algorithm; Colorado includes it; Montana excludes "downstream physical effects."[17] These are not pedantic differences. A law that protects only the raw signal and not the inference protects the lock but not the valuables. Any legislator drafting here should treat the *definition* as the single most consequential clause in the bill.

The instruments side by side

Jurisdiction	Instrument	Force	What it covers
UNESCO (global)	Recommendation on the Ethics of Neurotechnology, 2025	soft law	Full life cycle; values + policy actions; consent, mental privacy, enhancement limits
Chile	Constitutional amendment (Law 21.383, 2021) + case law	binding	Brain activity & derived information; psychic integrity
Spain	Charter of Digital Rights, 2021	guiding	Neurorights principles in a digital-rights context
Brazil	Constitutional amendment on mental integrity	binding	Mental integrity as a protected interest
EU	GDPR; AI Act; (OECD 2019)	indirect	Data protection + narrow AI prohibitions; no dedicated instrument
US states	CO, CA (2024); MT, CT (2025); ~15 pending	binding	Neural data as "sensitive"; consent / deletion rights
US federal	MIND Act (proposed, 2025)	pending	Directs FTC study; no substantive protection yet

4. Sovereignty of Thought — Guaranteeing the inner forum

Beneath all five neurorights lies an older one: freedom of thought, protected in Article 18 of the International Covenant on Civil and Political Rights. It has always been treated as *absolute* — one of the very few rights admitting no exception, no derogation even in war. The reason it could be absolute is that it was unenforceable from the outside: no state could actually reach into a head to police a thought, so the right cost nothing to grant in full.

Neurotechnology removes that comfortable impossibility. For the first time, the *forum internum* — the inner forum of the mind — is becoming technically accessible. "Sovereignty of thought," then, means something concrete: a guarantee that the inside of your head is the one domain where you remain the only authority, where no employer, platform, insurer or government may read, infer, or alter without your free and specific consent, and where some interventions are forbidden even *with* consent, because they would erode the conditions for consent itself.

In practice, a real guarantee requires four things at once. **First, a default of inaccessibility**: neural data should be opt-in only, never the price of using a device. **Second, purpose limitation with teeth**: data gathered to help you sleep cannot be repurposed to score you for a loan. **Third, prohibitions that survive consent**: just as a person cannot validly consent to be enslaved, there should be neural interventions — covert influence, manipulation that bypasses conscious deliberation — that no terms-of-service checkbox can authorise. **Fourth, enforceability**: rights without remedies are decoration. The Chilean case mattered precisely because a court ordered data deleted.

5. Ownership — Can a company own your thoughts?

Under most current law, the uncomfortable answer is: closer to "yes" than anyone is comfortable admitting. When you accept the terms of service of a consumer EEG device, you typically grant the manufacturer a broad licence to store, process and often share the data your brain generates. A 2024 report by the Neurorights Foundation examined thirty consumer neurotechnology companies and found that twenty-nine of the thirty had access to users' brain data and imposed essentially no meaningful limits on that access, with nearly all reserving the right to share it with third parties.[14][18] Your thoughts, rendered as data, are being treated like any other exhaust of digital life — to be licensed, brokered and modelled.

The deeper problem is that property law is a poor fit for thought. As legal scholars analysing the UNESCO Recommendation have noted, neural data and the probabilistic inferences drawn from it form a genuinely novel class, and the familiar tools map badly onto something so intimate.[19] Three broad approaches are in contention:

- **Property / ownership models** grant you a property right in your neural data, which you can then license or sell. Intuitive but dangerous: anything you can sell, you can be pressured or impoverished into selling, and a market in thoughts will find the desperate first.
- **Data-protection models** (the EU and US-state approach) treat neural data as an especially sensitive category, with heightened consent and security duties. The path of least resistance, but it inherits all the weaknesses of consent-based privacy: people click "agree" without reading, and the protection collapses.
- **Bodily-integrity / inalienable models**, favoured by the Morningside Group, treat neural data like a transplanted organ — something that may be *donated* for altruistic or medical purposes but never *commercialised*, regardless of consent.[3][20] The strongest protection, and the hardest to enact, because it forecloses a business model investors are counting on.

For a legislator, the question should be reframed: *which* model do we choose, and do we choose it before or after the markets are built around the weakest one? The window in which that choice is cheap is closing.

6. The Right Level — At what level should neuro-rights be legislated?

Neural data crosses borders at the speed of a cloud upload; a Chilean's brain data was processed on servers in San Francisco. Protection at any single level is therefore necessary but never sufficient. The realistic answer is that neurorights must be legislated at every level at once, with each layer doing the job it is uniquely suited to.

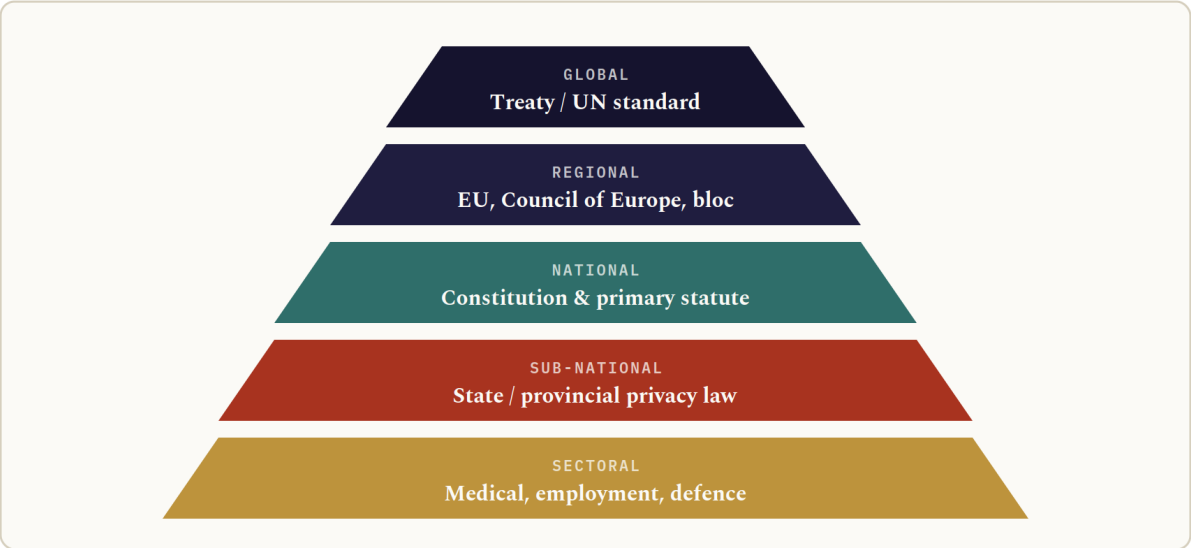


Figure 3. No single tier is enough. Global sets norms and can ban the worst uses; national supplies the foundational right and enforcement; sectoral rules bind the doctors, employers and soldiers who actually deploy the technology.

My own view, offered as analysis rather than dogma: the **foundational protection belongs at the national constitutional level**, because only a constitution can make a right resistant to the ordinary political bargaining that erodes statutes. Chile shows it can be done. But constitutional text alone is too abstract to bind a product manager, so it must be paired with **detailed sectoral regulation** — the granular rules for clinical devices, workplace monitoring and the rest — and with a **global instrument with actual teeth** for the cross-border and weapons questions no nation can solve alone. The UNESCO Recommendation is the seed of that global layer; the task of the next decade is to grow it from a recommendation into a convention.

7. The Gaps — What key legislation is missing?

If you lay the existing instruments side by side, the holes are as informative as the coverage. The most serious gaps, in rough order of urgency:

- **Inferred mental states.** Most laws protect the neural *signal* and ignore the algorithmic *inference* — the very thing that reveals your mood, attention or political reaction. Closing this gap is the single highest-value fix available.
- **Non-consumer use.** Existing statutes mostly target companies selling to consumers. They are largely silent on the state, the employer, the school and the insurer — the actors with the most power and the strongest incentive to read minds without genuine consent.

- **Manipulation and output, not just input.** Nearly all law concerns reading the brain. Almost none addresses *writing* to it — stimulation, closed-loop systems that adjust your state, and the manipulation risk that follows.
- **Redress.** There is no established cause of action for "my mind was read or altered without consent," no settled measure of damages, no specialist forum.
- **Detection and audit standards.** No jurisdiction yet requires neurotech systems to keep tamper-evident logs of what they read and wrote, which would make abuse provable after the fact.
- **Weaponisation.** No binding international instrument prohibits the development or use of neuroweapons. This is the largest hole of all, and the subject of the next section.

8. The Security Frontier — Neuroweapons, cognitive warfare, and the front line

This is the part of the subject most prone to both reckless hype and reckless dismissal, so it deserves the most careful handling. Two things are simultaneously true: military and intelligence establishments treat the brain as an emerging operational domain; and the most dramatic public claims about deployed mind-control weapons remain unproven and, in important cases, have been formally judged unlikely. A serious policy must hold both facts at once.

The doctrine is real

NATO has formally explored "cognitive warfare" as an emerging domain, and its associated studies have used arresting language — describing the human brain as a contested battlefield of the twenty-first century.[21][22] This is doctrine, not data: it describes how strategists are *thinking*, and it spans a wide spectrum, most of which is non-invasive. At the soft end, cognitive warfare means disinformation and algorithmic manipulation of attention and emotion through social media — influencing minds at scale without touching a single neuron. At the hard end sit hypothesised directed-energy or acoustic devices intended to physically disrupt the nervous system. The soft end is unquestionably happening now. The hard end is where the evidence becomes genuinely disputed.

The Havana Syndrome evidence, told straight

Beginning in 2016, US diplomats and intelligence officers reported a cluster of symptoms (pressure, pain, dizziness, cognitive fog) that became known as Havana Syndrome, or Anomalous Health Incidents. A 2020 US National Academies report concluded that "directed, pulsed radiofrequency energy" was the most plausible mechanism for a *subset* of cases. That finding is frequently cited as proof of an energy weapon. But the story did not end there, and intellectual honesty requires the rest of it:

> **The contested state of the evidence.** In 2023, a US intelligence-community assessment concluded it was *very unlikely* that a foreign adversary using a weapon was responsible, attributing many incidents to pre-existing conditions, environmental factors and other causes. Large imaging studies published in 2024 found no significant differences between the brains of affected personnel and matched controls —

> contradicting an earlier, smaller study. The matter remains politically live. The responsible position is not "it was a weapon" nor "it was nothing," but: **the dramatic mechanism is unproven, official assessments have leaned against it, and the question is still open.** Policy should be robust whether or not such weapons exist — not built on the assumption that they certainly do.[23][24]

Where do you mark the front line?

The hardest conceptual problem in regulating neuroweapons is that the front line is invisible and may run through a civilian's living room. Conventional weapons announce themselves; a cognitive weapon, by design, would not. This is what distinguishes the two postures a state might adopt:



Figure 4. The cold-war column is the more dangerous one for civil liberties, precisely because it operates below the threshold of declared conflict where the laws of war and most domestic protections do not reach.

Would non-consensual use ever be ethical?

Almost never — but the honest answer is not "never," and pretending otherwise weakens the argument. The defensible exceptions are narrow and familiar from other law: a person who lacks decision-making capacity and faces imminent, grave harm may be treated without contemporaneous consent (the logic that licenses emergency surgery on an unconscious patient); a court might, under strict due process, compel a narrow neural measurement much as it compels a blood draw. What is *not* defensible is covert influence, interrogation by forcible brain-reading, or any manipulation that bypasses a competent person's conscious will. The line that holds: measurement of a non-competent person to save them may be justified; alteration of a competent person against their will is not.

A global ban with an enforced moratorium?

There is a strong case for a binding international prohibition on neuroweapons designed to covertly impair or control cognition, modelled on the Biological and Chemical Weapons Conventions. In favour: such weapons would be uniquely deniable, would target the seat of human dignity itself, and — like biological weapons — offer little stable advantage once both sides possess them. Against: a ban is only as good as its verification, and you cannot inspect for a weapon you cannot detect. A moratorium that cannot be verified mainly disarms the law-abiding. The realistic path is therefore a ban *coupled with* a serious, funded programme to build the detection science verification would require — which is the next problem.

Detecting illegal use

Detection splits into two very different problems. Detecting **illegal data practices** — a company reading neural data without consent, an employer monitoring brains — is tractable today through ordinary tools: mandatory disclosure, audits, tamper-evident logging, and penalties severe enough to bite. Detecting an alleged **covert neuroweapon** is far harder and, candidly, beyond current capability for the more exotic claims. The constructive response is to invest in the science that *could* make it tractable: environmental RF and acoustic monitoring in sensitive facilities; standardised forensic protocols for anomalous health incidents; and a culture of measuring rather than assuming. The honest message to constituents who believe they are victims is that the field's *inability to detect* is itself a policy failure worth fixing — not evidence that their distress is unreal, and not, by itself, evidence of the mechanism they suspect.

9. The Evidentiary Problem — Could you prove, after the fact, that a mind was controlled?

This deserves a direct and unsentimental answer, because a great deal hangs on it. **With today's science, no — not for the strong claim of covert, remote mind control.** There is no validated forensic test that can examine a person and establish that their thoughts or actions were externally directed. This is not because the question is forbidden; it is because the capability the test would have to detect has not been demonstrated to exist, and the brain's normal variability is so large that distinguishing "influenced" from "ordinary" states is, at present, not possible to a forensic standard.

What forensic neuroscience *can* do is narrower and worth stating precisely. It can sometimes document physical injury consistent with energy exposure (as the Havana Syndrome imaging debate showed — inconclusively). It can establish whether an implanted device was present and what it logged, *if* the device kept honest logs. It can document a pattern of coercion or surveillance through ordinary evidence — records, witnesses, communications. What it cannot do is read the residue of an alleged covert influence off a brain scan.

> **Why this is an argument for regulation.** The provability gap is not a reason to dismiss the concern; it is a reason to build the infrastructure that would close it. If neurotech devices were legally required to maintain tamper-evident, independently auditable logs of every signal read and every stimulation delivered — the neural equivalent of an aircraft's black box — then abuse by a *known* device would become provable after the fact. That is an achievable legislative goal today, and the single most useful thing a lawmaker could do for anyone who fears their mind has been tampered with. We cannot yet detect the exotic; we can absolutely make the deployed accountable.

10. The Benevolent Case — What good can it do, beyond helping the disabled?

It would be a grave mistake to treat neurotechnology only as a threat. The same capability that frightens us is, in clinical hands, one of the most hopeful developments in modern medicine. The restoration cases are already real: people with paralysis controlling cursors, robotic arms and tablets through implanted interfaces; deep-brain stimulation calming the tremor of Parkinson's and, increasingly, treatment-resistant depression and epilepsy; cochlear implants and emerging visual prostheses; and laboratory systems that have decoded attempted or imagined speech to give a voice back to people who have lost theirs.[25][26] Five people with severe paralysis were reported using one company's implant by mid-2025; another's device has been guided into the brain through a blood vessel rather than open surgery.[26][27]

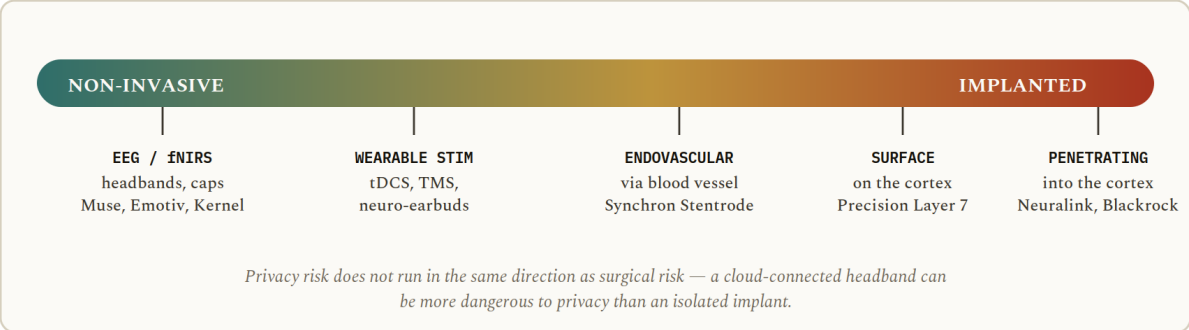


Figure 5. The current landscape, arranged by how deeply the technology reaches into the body. Privacy risk does not run in the same direction as surgical risk.

Beyond restoring lost function, the plausible benevolent horizon is broad: objective, continuous monitoring that could transform psychiatry from a discipline of self-report into one of measurement; early detection of stroke, seizure and neurodegeneration; closed-loop therapies that deliver stimulation only when symptoms appear; communication for the "locked-in"; safer attention monitoring for pilots and long-haul drivers; and, further out, genuine augmentation of memory and learning. The benefits are real enough that a blanket prohibition would be both impossible and cruel. The task is not to stop neurotechnology but to ensure its benefits arrive without its harms — which is exactly what rights frameworks are for.

11. The Multiplier — How AI changes everything about neurotech

It is impossible to reason about neurorights without reasoning about artificial intelligence, because AI is the component that turns a stream of brain signals into knowledge about a person. As one neurorights advocate put it, the usability and significance of brain data is now driven by pattern recognition; the device captures the signal, but the model finds the meaning.[13] This changes the risk calculus in several specific ways.

- **Decoding leaps ahead of consent.** A device you authorised to track your sleep may, with a better model applied to the same raw data later, suddenly reveal far more

than you agreed to share. The consent you gave was for what the data could reveal *then*, not what it can reveal now.

- **Foundation models for the brain.** Like language models, "brain foundation models" trained on large neural datasets promise to decode new individuals with little calibration — which means your brain data, pooled with millions of others, makes everyone easier to read, including people who never consented.
- **Closed loops invite manipulation.** When an AI both reads your state and decides what stimulation or content to feed back, the line between therapy and manipulation becomes a matter of who sets the objective function.
- **Inference over collection.** The danger migrates from *what is measured* to *what is inferred* — precisely the category most laws fail to protect.

The policy implication is blunt: neurorights law that does not govern the AI applied to neural data is protecting a locked door beside an open window.

12. The Trade — What rights will we hand over, and have we already?

If the last twenty years of digital life teach anything, it is that people trade intimate data for convenience readily, repeatedly, and without much regret — until the consequences arrive. We surrendered location, contacts, browsing and biometrics for free services and faster phones. There is no reason to expect neural data will be different unless the rules are different. The likely sequence is familiar: a genuinely useful application appears; using it requires sharing neural data; the sharing becomes normal; the data becomes an asset; and the asset is used in ways the original users never imagined — pricing, hiring, advertising, scoring.

The rights most at risk of quiet surrender are mental privacy (traded for personalisation) and a slice of agency (traded for the convenience of a system that anticipates your wants). The decisive question for a legislator is whether the most sensitive category of data humanity has ever generated will be governed by the same click-to-accept consent regime that has so comprehensively failed everywhere else — or whether some neural rights will be made *inalienable*, so they cannot be signed away no matter how tempting the app.

13. Society — How everyday life changes

The second-order effects reach well beyond privacy. In the **workplace**, neural monitoring of attention and fatigue could improve safety and just as easily become the most invasive surveillance ever deployed against employees. In **education**, attention-tracking promises personalised learning and threatens to discipline children's inner states. In the **justice system**, the temptation to use neural data as a lie detector or risk score will be enormous and, given the science, mostly unwarranted. In the **economy**, a "neuro-divide" looms: if cognitive augmentation works and is expensive, it could entrench inequality at the level of the brain itself — which is why fair access to augmentation is one of the five rights. And in **culture**, a society in which inner states are routinely measured may slowly lose the very idea of a private

inner life — the assumption, woven through every legal and moral tradition we have, that what you think is yours until you choose to share it.

14. Victims — Defining, recognising and helping them

"Victim of neurotechnology" needs a careful definition, both to help real victims and to avoid a category so loose it helps no one. The defensible definition centres on *demonstrable harm arising from the use or misuse of a neurotechnology*, which falls into recognisable classes:

- **Data-rights victims** — people whose neural data was collected, inferred from, shared or sold without valid consent. The largest and most provable class, with the Chilean case as the template for redress.
- **Clinical and research harm** — patients or trial participants injured by a device, an explantation, a malfunction, or inadequate consent.
- **Coercion and surveillance victims** — people pressured by an employer, institution or state into neural monitoring, or monitored covertly.
- **Discrimination victims** — people disadvantaged in hiring, insurance, credit or justice on the basis of neural data or inferences.

> **A note written with care.** A growing number of people report believing they are victims of covert, remote neural influence. Their distress is real and deserves a serious, humane response — not reflexive dismissal. At the same time, an honest framework has to acknowledge what Section 9 establishes: there is currently no validated way to confirm covert remote influence, and for the specific symptoms most often described, mainstream medical assessment finds a range of explanations, many treatable. The compassionate *and* truthful response is the one good medicine already offers: take the suffering seriously, pursue thorough clinical evaluation, document everything, and press for the detection science and device-accountability rules that would let any future claim be tested rather than merely believed or denied. Helping victims means building systems that can tell the difference — not assuming the answer in either direction.

15. The Horizon — 3, 5, 10 and 20 years

Forecasts in this field age badly, so treat the following as scenario-planning, with confidence decreasing sharply as the horizon lengthens. The market is real — well over a billion dollars in fresh investment flowed into brain-computer interfaces in 2025–26 alone — but no implanted BCI has yet been approved for general medical use; all remain experimental.[27][28]

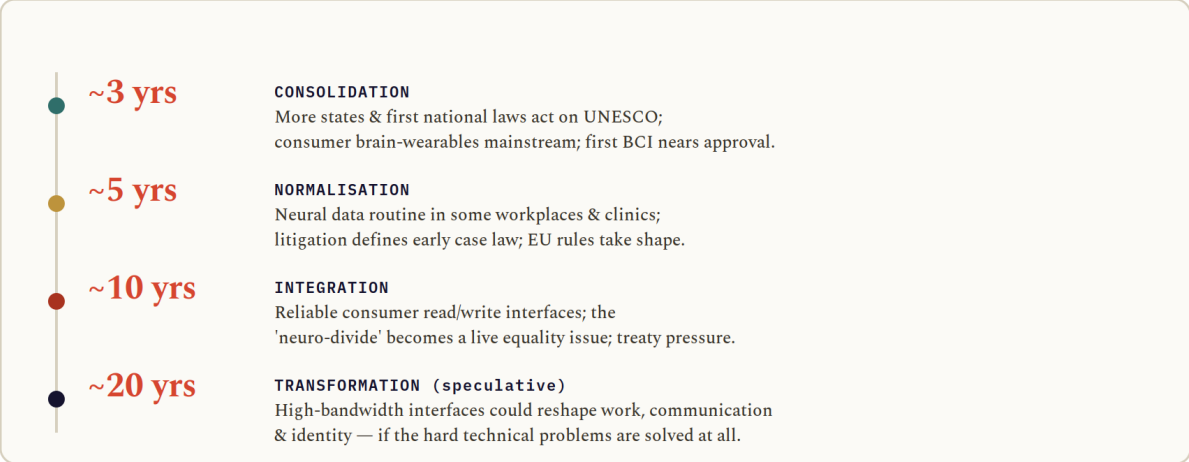


Figure 6. Near-term forecasts concern law and markets and are fairly reliable; long-term forecasts concern capability and are not. The honest twenty-year statement is a conditional, not a prediction.

16. Follow the Money — Who benefits commercially, and how

Understanding the incentives is half of understanding the politics. The value does not sit only with the device makers everyone names.

- **Device makers** — the implant and wearable companies (Neuralink, Synchron, Blackrock Neurotech, Precision Neuroscience, Paradromics, Kernel, Emotiv, Muse). They monetise hardware, and increasingly the data and software around it.[26][29]
- **Platforms and big tech** — whoever owns the operating layer that brain interfaces plug into captures enormous value; a major consumer-electronics firm moving to define a standard input protocol for BCIs in 2025 is a signal of where this is heading.[30]
- **Data brokers and advertisers** — the most lucrative and least visible beneficiaries, for whom neural inference is the richest targeting signal ever produced.
- **Employers and insurers** — who benefit from knowing the attention, stress and health states of workers and policyholders.
- **Pharma and clinical providers** — through closed-loop therapeutics and objective biomarkers.
- **Defence and intelligence** — the least transparent sector, and the reason so much capability sits behind classification.

The pattern to watch is the one that defined the social-media era: the company that sells you the helpful device is rarely the one that profits most from the data it generates. Regulation

aimed only at device makers will miss most of the money — and therefore most of the incentive to misbehave.

17. The Classified Question — What might exist behind secrecy, and why

Any honest treatment must address what is secret, while refusing to pretend that secrecy is the same as evidence. The responsible method is to separate the documented record from speculation, and to label each.

What is documented. The historical record contains real, declassified programmes that explored manipulating the mind: the United States' MKUltra (mind-control and interrogation research, 1950s–70s, largely conducted without consent and later condemned), and Cold-War-era research on microwave and behavioural effects on both sides. These are not conspiracy theories; they are matters of public record, exposed by official inquiry. They establish a crucial point for policy — that states have, in living memory, conducted non-consensual neural experimentation — which is reason enough for binding oversight regardless of what exists today.

What is plausibly classified, and why. It is reasonable to assume several states fund classified neuroscience with defensive and offensive aims, because the public defence-research record implies a classified counterpart, as it does in every dual-use field. Classification here is driven by ordinary, non-sinister logic: protecting a capability advantage, protecting sources and methods, and avoiding the diplomatic cost of confirming research into prohibited-seeming areas. None of that requires — or supports — the strong claim that deployed devices are covertly controlling civilian populations.

> **Where speculation ends.** Speculating that classified *research* exists is well grounded. Asserting that a fielded, covert, civilian-targeting mind-control weapon exists is not — it remains unproven, and the most exotic public claims have, where formally examined, been judged unlikely (Section 8). The intellectually honest stance for a policymaker is to legislate for the *capability that history proves states will pursue*, build the detection and oversight that would catch abuse, and decline to either dismiss or credit unverified claims. Good policy does not require knowing the secret; it requires assuming states keep some, and binding them anyway.

18. Action — Immediate steps, by institution

If this essay has an operational core, it is the table below. Each branch and sector has something it can do now, without waiting for the others.

Institution	Immediate action
Legislative branch	Define "neural data" to include *inferences*; classify it as a special protected category; make a core of neural rights *inalienable* (non-waivable by consent); fund a regulator. Constitutionalise the foundational right where possible.
Executive branch	Adopt the UNESCO Recommendation domestically; convene a national neurotechnology strategy across health, labour, justice and defence; fund detection and forensic science; require neural-device "black-box" logging in procurement.
Judiciary	Prepare doctrine now: recognise neural data within existing privacy, bodily-integrity and dignity jurisprudence; set evidentiary standards that exclude neural "lie detection" until validated; create capacity to hear redress claims.
Law enforcement	Treat covert neural data collection as a serious offence; build technical capacity to investigate device-based abuse; resist adopting neural surveillance tools without statutory authority and oversight.
Medical sector	Tighten informed-consent standards for neural devices and data; insist neural data donated for research is not commercialised; develop standardised forensic protocols for anomalous neurological incidents; lead, rather than follow, on regulation.

19. Oversight — What it should look like

Oversight that exists only on paper is worse than none, because it manufactures false confidence. Effective oversight of neurotechnology needs five components, layered like the governance itself:

- **A competent regulator** with technical staff, the power to license high-risk neural products before market, and the power to fine and ban — not a committee that issues guidance.
- **Mandatory pre-deployment assessment** — an ethical and rights impact assessment for any system that reads or writes to the nervous system, of the kind UNESCO's implementation tools envisage.[6]
- **Auditable transparency** — the black-box logging requirement, so that what a device does to a brain is recorded and independently inspectable.
- **Independent red-teaming** — funded adversarial testing of devices and the AI models behind them, before and after release.
- **An international node** — a body to handle cross-border and weapons questions, with eventual verification powers; the seed of the binding treaty that soft law should grow into.

The guiding principle is proportionality: a meditation headband and a cortical implant should not face identical rules, but both should face *some* rule, scaled to how deeply they reach and how widely their data travels.

In Closing — The window, and who holds it open

The history of technology regulation is a history of arriving late. We wrote the rules for the automobile after the deaths, for finance after the crashes, for social media after the damage was done. Neurotechnology offers something rare: a chance to legislate *before* the harm becomes structural, while the industry is small enough to shape and the public is still alarmed enough to act. UNESCO has supplied a shared vocabulary and a set of values. Chile has shown that constitutional protection is possible and enforceable. A handful of US states have shown that even partial protection can be passed quickly. What is missing is not the idea but the will — the decision, by enough legislators in enough places, to treat the inside of a human head as the one place that markets and states may not enter without permission.

The mind has been private for as long as there have been minds. Whether it stays that way is now, for the first time, a policy choice. It will be made in the next handful of years, by people who can either read this subject seriously or wait until it is too late to matter. This essay is an argument for reading it seriously, and for choosing — deliberately, and soon — to keep the last private place private.

A note on method: where the evidence is solid — the laws, the rulings, the technology that exists — this essay states it plainly and cites it. Where the evidence is contested or absent — the exotic weapons, the classified capabilities, the strongest claims of covert control — it says so, and resists the temptation to resolve uncertainty in either direction for effect. On a subject that touches the deepest fears and hopes we have about our own minds, that discipline is not a hedge. It is the whole point.

References

1. UNESCO. “Ethics of neurotechnology: UNESCO adopts the first global standard.” unesco.org
2. UNESCO. “Towards an International Instrument” & Recommendation page. unesco.org/en/ethics-neurotech/recommendation
3. Yuste, R. et al. “Four ethical priorities for neurotechnologies and AI.” Nature (2017); Morningside Group / NeuroRights Foundation.
4. Columbia University, Institute for the Study of Human Rights. “Neurorights: Human Rights Guidelines for Neurotechnology and AI.”
5. Borbón, D. & Borbón, L. “A Critical Perspective on NeuroRights.” Frontiers in Human Neuroscience (2021).
6. Covington / Inside Privacy; Global Policy Watch. “UNESCO Adopts First Global Framework on Neurotechnology Ethics” (2026).
7. ArentFox Schiff. “Neural Rights: Landmark Ruling” (Chile, Law 21.383).
8. Mantegna, M. “Neural Data on Trial: Chile’s Supreme Court addresses the First Global Neurorights Case.” Medium (2023).
9. Future of Privacy Forum. “Privacy and the Rise of ‘Neurorights’ in Latin America.”
10. “Chilean Supreme Court ruling on the protection of brain activity” (rol N°1.080-2020). PMC10929545.
11. Stanford Law School, Law & Biosciences Blog. “Even Chile’s Neurorights Leave Inferred Mental Data in a Gray Zone” (2026).
12. “Is Your Neural Data Part of Your Mind?” / “Is Mental Privacy a Component of Personal Identity?” PMC8460199 & PMC8551354.
13. KFF Health News. “States Pass Privacy Laws To Protect Brain Data Collected by Devices” (2025).
14. CBS News / Colorado Public Radio. Reporting on CO, CA, MT neural-data laws; AMA & FTC actions (2025).
15. Cooley LLP. “Wave of State Legislation Targets Mental Privacy and Neural Data” (2025).
16. Davis Wright Tremaine. “U.S. Senators Propose ‘MIND Act’…” (2025).
17. Future of Privacy Forum. “The ‘Neural Data’ Goldilocks Problem” (2025).
18. Genser, J., Damianos, S. & Yuste, R. “Safeguarding Brain Data” (Neurorights Foundation report on 30 consumer companies).
19. Stanford Law School. “Who Owns Digital Thoughts? The Limits of Property Law and the 2025 UNESCO Recommendation” (2026).
20. Goering, S. et al. “Recommendations for Responsible Development and Application of Neurotechnologies.” Neuroethics (2021).
21. NATO Allied Command Transformation. Cognitive Warfare concept & exploratory studies. act.nato.int
22. “Neuro-Cognitive Warfare: Inflicting Strategic Impact via Non-Kinetic Threat.” Small Wars Journal (2022).
23. U.S. National Academies. “An Assessment of Illness in U.S. Government Employees and Their Families at Overseas Embassies” (2020).
24. U.S. ODNI / Intelligence Community. Assessment on Anomalous Health Incidents (2023); NIH/JAMA imaging studies (2024).
25. “Neurochips: The State of Brain-Computer Interfaces in 2025.” andersenlab.com
26. Clinical Trials Arena. “Brain-computer interfaces are closer than you think” (2025).
27. bciintel.com. “State of BCI: 2026 Annual Industry Report” (funding & market data).
28. Mordor Intelligence. “Neurotechnology / Brain-Computer Interface Market” (2026).
29. Ambula Healthcare / Spherical Insights. Company profiles: Neuralink, Synchron, Blackrock, Precision Neuroscience, Kernel (2025).
30. Apple BCI Human Interface Device input protocol (2025); Synchron iPad demonstration.
31. Ienca, M. & Andorno, R. “Towards new human rights in the age of neuroscience and neurotechnology.” Life Sciences, Society and Policy (2017).
32. OECD. “Recommendation on Responsible Innovation in Neurotechnology” (2019).