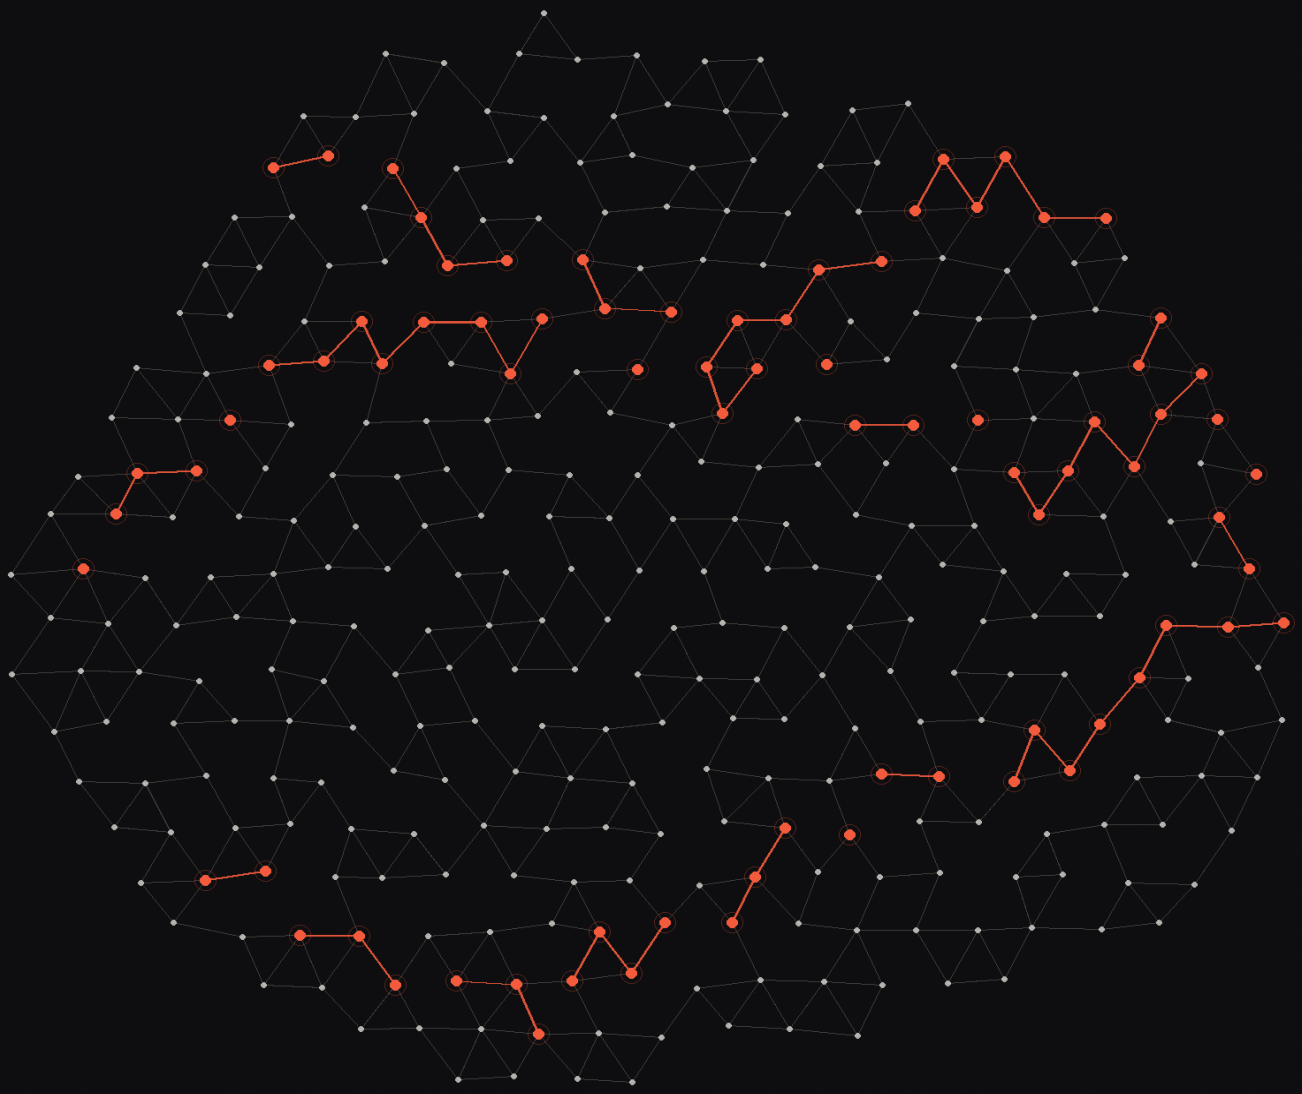
FIG. 01 / CORTICAL SIGNAL / NON-CONSENSUAL

47.4°N · 4.3°E



NEURO— RIGHTS.

The Battle for the Last Private Place

STEPHANE VAN DER AA

TI ONE VOICE
one.witysk.org

VOL. IV / MMXXVI § ISBN — / PRINTED 2026

NEURO-RIGHTS

The Battle for the Last Private Place

A Comprehensive Guide for Citizens,
Legislators, and the Global Community

EXPANDED FOURTH EDITION

2026

Version 4.1



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.

PREFACE TO THE EXPANDED EDITION

This expanded fourth edition of *Neuro-rights: The Battle for the Last Private Place* responds to the rapid evolution of the field since the first edition. Four developments have necessitated significant new material.

This fourth edition (2026) adds substantial new material in response to developments since the third edition: a consolidated reference to existing laws at every level (Chapter 18-bis); a comprehensive device-maker atlas (Chapter 62-bis); a dedicated treatment of the vectors and consequences of mind-control risk (Chapter 33-bis); a complete set of model legislative provisions at all six tiers, from supra-national treaty to municipal rule (Chapter 67-bis); an institutional architecture for implementation (Chapter 70-bis); a geopolitical analysis over four time horizons (Chapter 76-bis); a dedicated chapter on synthetic telepathy and the conditions under which it must be constrained (Chapter 78-bis); and a new Part XV-bis containing fifteen fictional hypothetical cases, ranging from civilian to warfare to covert to future-technology contexts, each followed by legal commentary distinguishing what existing law can already do from what new doctrine or legislation would be required. New diagrams accompany the new chapters.

First, the scientific foundations of neurotechnology have advanced faster than the previous edition anticipated. Foundation-model approaches to neural decoding, mature consumer brain-computer interface products, and the integration of artificial intelligence with neural data collection have changed the technical landscape. This edition adds a primer on the neuroscience underlying these technologies, an analysis of foundation-model decoders, and an updated assessment of capability trajectories.

Second, the international legal landscape has expanded substantially. The UNESCO Recommendation on the Ethics of Neurotechnology, adopted in November 2025, established the first global ethical framework. State-level legislation in the United States, regulatory developments in the European Union, and constitutional initiatives in several Latin American countries have created a more complex international picture. This edition treats these developments in detail and adds comparative analysis across jurisdictions.

Third, the Targeted Individual phenomenon — a body of human testimony reporting non-consensual technological targeting — has reached a scale and consistency that demands serious engagement. The previous edition treated this material in passing; this edition devotes a full Part to it, examining the documented patterns, the psychological impact on affected individuals, and the diagnostic and forensic challenges that the phenomenon raises. The treatment is neither credulous nor dismissive: it takes seriously a body of testimony that current institutions have failed to address adequately, while maintaining critical distance from any particular causal theory.

Fourth, the need for practical guidance has become clear. Readers of the previous edition asked: what can I do? This edition adds a Practical Toolkit for citizens, including specific guidance on personal information hygiene, civic engagement, and the construction of local networks of solidarity and mutual aid. The work of defending neuro-rights, this edition argues, is the work of ordinary citizens at the scale of daily life — not only of legislators, scientists, and activists.

The structure of the book has been reorganized to support these additions while preserving the material from the previous edition. The original fifteen parts have been expanded to seventeen, with new content interleaved into existing parts and with two entirely new parts: Part V on the Targeted Individual phenomenon and Part XVI on the Practical Toolkit. Existing chapters have been retained largely in their original form, with light editing for currency and consistency.

The book is not a balanced presentation. It is an argument: that the human mind is a frontier worth defending, that the technologies challenging that frontier require legal and ethical infrastructure that does not yet exist, and that the construction of that infrastructure is among the most consequential civic projects of our time. Readers will encounter material that may be unfamiliar, contested, or unsettling. The goal is not comfort but engagement with what is, on any honest assessment, a defining issue of the present century.

A note on sources: the book draws on a wide range of materials, including UN reports, academic literature, declassified government documents, technical research papers, journalistic investigation, and the testimony of affected individuals. The treatment of contested claims is explicit: where evidence is strong, the text says so; where claims remain contested, the text acknowledges the contestation. The reader is invited to consult the cited sources and form independent judgments.

Finally, a note on the future of the project. This is the fourth edition; further revisions are anticipated as the field continues to evolve. Readers are invited to engage with the broader community working on these issues: at one.witysk.org, through national neuro-rights organisations, and through the local networks whose construction is among the recommendations of this book. The project of defending the last private place is collective; this book is one contribution among many.

CONTENTS

Open this contents list in Microsoft Word and press F9 (or right-click → Update Field) to populate page numbers.

CHAPTER 1 13
What Is Neurotechnology? A Layperson's Guide 13
CHAPTER 2 17
The Spectrum of Neurotechnologies — From EEG Headbands to Brain Implants 17
CHAPTER 3 20
How the Brain Actually Works — A Neuroscience Primer 20
CHAPTER 4 25
Reading and Writing the Brain — Decoding and Encoding 25
CHAPTER 5 29
What Are Neuro-rights? The Five Core Rights Explained 29
CHAPTER 6 32
Cognitive Liberty — The Foundational Freedom 32
CHAPTER 7 35
Why Now? The Urgency of Neuro-rights in the 2020s 35
CHAPTER 8 41
Chile's Pioneering Constitutional Reform (2021) 41
CHAPTER 9 43
The UN Human Rights Council Resolution 51/3 (2022) 43
CHAPTER 10 45
The 2024 UN Report on Neurotechnology and Human Rights 45
CHAPTER 11 49
The 2025 UNESCO Recommendation on the Ethics of Neurotechnology 49
CHAPTER 12 52
Spain, Brazil, and the Latin American Neuro-rights Movement 52
CHAPTER 13 57
China and the Authoritarian Use of Neurotechnology 57
CHAPTER 14 60
The African Union, the Arab World, and Global South Perspectives 60
CHAPTER 15 64
Colorado and California — US State-Level Neurodata Protection 64
CHAPTER 16 67
The EU AI Act and Its Gaps Regarding Neurotechnology 67
CHAPTER 17 70
The Council of Europe Framework Convention on AI 70
CHAPTER 18 71
Comparative Analysis — A Patchwork of Protections 71
CHAPTER 18-bis 75
Existing Laws at Every Level — A Consolidated Reference 75
CHAPTER 19 83
The Forum Internum — Why Your Thoughts Are Sacred 83
CHAPTER 20 86
The Philosophy of Mental Privacy — Locke to Bostrom 86
CHAPTER 21 90
Mental Privacy in the Age of Brain-Reading 90
CHAPTER 22 94

Intellectual Property Rights to Thoughts — Can Companies Own Your Brain Data? 94
CHAPTER 23 97
Neurodiscrimination — When Your Brain Activity Determines Your Future 97
CHAPTER 24 103
The History of Cognitive Warfare — From Propaganda to Neuroweapons 103
CHAPTER 25 107
Cognitive Warfare — The Battle for the Brain 107
CHAPTER 26 111
Havana Syndrome — The First Salvo of Neurowar? 111
CHAPTER 27 114
DARPA's N3 Program — Non-Surgical Neurotechnology for the Military 114
CHAPTER 28 117
Russia's Psychotronic Programmes — A Historical Survey 117
CHAPTER 29 121
China's Military Brain Science Programme 121
CHAPTER 30 125
NATO and the Cognitive Domain 125
CHAPTER 31 128
The Neuroweapons Arsenal — What Exists and What Is Possible 128
CHAPTER 32 132
Hot War vs Cold War — How Neuroweapons Change Conflict 132
CHAPTER 33 135
Where Is the Frontline? Defining the Battlefield of the Mind 135
CHAPTER 33-bis 138
Where Mind-Control Risk Comes From — Vectors, Actors, Consequences 138
CHAPTER 34 144
The Targeted Individual Phenomenon — An Introduction 144
CHAPTER 35 147
Documented Patterns in Targeted Individual Testimony 147
CHAPTER 36 151
The Psychological Impact of Targeted Individual Experiences 151
CHAPTER 37 154
Differential Diagnosis — Technological Targeting Versus Psychiatric Illness 154
CHAPTER 38 163
The Case for a Global Ban on Neuroweapons 163
CHAPTER 39 168
An Enforced Moratorium — How Would It Work? 168
CHAPTER 40 171
How to Detect Illegal Neurotechnology Use 171
CHAPTER 41 175
Proving Mind Control After the Fact — Forensic Challenges 175
CHAPTER 42 179
Defining a Victim of Neurotechnology — Recognition and Support 179
CHAPTER 43 185
The Absolute Prohibition — Torture and Neurotechnology 185
CHAPTER 44 188
Medical Exceptions — When Is Non-Consensual Use Ethically Justified? 188
CHAPTER 45 191
Criminal Justice — Brain-Reading as Evidence 191
CHAPTER 46 193
The Slippery Slope — From Consent to Coercion 193
CHAPTER 47 199
Medical Miracles — Treating Parkinson's, Epilepsy, Depression, and Paralysis 199
CHAPTER 48 202
Cognitive Enhancement — Smarter, Faster, Better? 202

CHAPTER 49 205
Neuroeducation — Learning at the Speed of Thought 205
CHAPTER 50 208
Mental Health — New Frontiers in Treatment 208
CHAPTER 51 213
Human-Machine Teaming — The Centaur Model 213
CHAPTER 52 216
AI as the Brain's Operating System 216
CHAPTER 53 219
Foundation Models for the Brain — The Coming Neural Decoder Revolution 219
CHAPTER 54 222
Algorithmic Bias in Neurotechnology 222
CHAPTER 55 225
AI and Cognitive Liberty — The New Control Challenge 225
CHAPTER 56 231
The Neuronal Divide — A New Class of Cognitive Haves and Have-Nots 231
CHAPTER 57 234
What Rights Are We Willing to Give Up? 234
CHAPTER 58 237
Neurotechnology and Democracy — The Threat to Free Elections 237
CHAPTER 59 240
Children and Neurotechnology — A Vulnerable Generation 240
CHAPTER 60 243
Workplace Neurosurveillance — The Coming Cognitive Panopticon 243
CHAPTER 61 249
The Neurotechnology Market — $7.3 Billion and Growing 249
CHAPTER 62 252
Big Tech's Brain Rush — Neuralink, Synchron, and the Silicon Valley Players 252
CHAPTER 62-bis 255
The Device Makers Atlas — A Complete Industry Overview 255
CHAPTER 63 260
The Military-Industrial-Neuro Complex 260
CHAPTER 64 263
The Investment Ecosystem — Who Funds the Future of the Brain? 263
CHAPTER 65 271
The Missing Legal Framework — Five Critical Gaps 271
CHAPTER 66 275
At What Level Should Neuro-rights Be Legislated? 275
CHAPTER 67 279
Model Legislation — Drafting the Neuro-rights Statute 279
CHAPTER 67-bis 284
The Ideal Legislative Stack — Complete Model Statutes at Every Level 284
CHAPTER 68 290
The Architecture of an International Neurotechnology Treaty 290
CHAPTER 69 294
Immediate Action Required — A Call to Five Sectors 294
CHAPTER 70 302
The Shape of Oversight — Independent Neurotechnology Commissions 302
CHAPTER 70-bis 307
Implementation — Existing Agencies and the New Agencies Needed 307
CHAPTER 71 311
Technical Audit Mechanisms — How to Verify Compliance 312
CHAPTER 72 316
Whistleblower Protection in the Neurotechnology Sector 316
CHAPTER 73 323

3 Years From Now (2029) — The Tipping Point 323
C H A P T E R 7 4 326
5 Years From Now (2031) — The Regulatory Race 326
C H A P T E R 7 5 329
10 Years From Now (2036) — The Neuronal Divide 329
C H A P T E R 7 6 332
20 Years From Now (2046) — The Transhuman Era 332
CHAPTER 76-bis 335
Geopolitical Impact — How Neurotechnology Reshapes World Politics 335
C H A P T E R 7 7 344
From MKULTRA to the Present — A History of Secret Brain Research 344
C H A P T E R 7 8 347
Speculative Technologies — What Likely Exists but Is Classified 347
CHAPTER 78-bis 351
Synthetic Telepathy — When It Will Exist, How It Must Be Constrained 351
CHAPTER XV-bis-1 355
Civilian Case One — The Employer Wearable (2027) 355
CHAPTER XV-bis-2 356
Civilian Case Two — The Schoolchild EEG (2028) 356
CHAPTER XV-bis-3 357
Civilian Case Three — Insurance Redlining (2030) 357
CHAPTER XV-bis-4 358
Civilian Case Four — The Smart-Home Decoder (2031) 358
CHAPTER XV-bis-5 358
Warfare Case One — Cognitive Psyop on Civilians (2028) 359
CHAPTER XV-bis-6 359
Warfare Case Two — Drone-Pilot Brain-Computer Teaming (2032) 359
CHAPTER XV-bis-7 360
Warfare Case Three — Targeted RF on Diplomats (2029) 360
CHAPTER XV-bis-8 361
Covert Case One — Activist Surveillance (2027) 361
CHAPTER XV-bis-9 362
Covert Case Two — The Untraceable Behavioural Nudge (2030) 362
CHAPTER XV-bis-10 363
Covert Case Three — Minority Tracking (2031) 363
CHAPTER XV-bis-11 364
Covert Case Four — Foreign Influence Campaign (2032) 364
CHAPTER XV-bis-12 364
Future-Tech Case One — The Memory Rewrite Trial (2034) 364
CHAPTER XV-bis-13 365
Future-Tech Case Two — The Synthetic-Telepathy Abuse (2036) 365
CHAPTER XV-bis-14 366
Future-Tech Case Three — The Neural Cloud Breach (2033) 366
CHAPTER XV-bis-15 367
Future-Tech Case Four — The Trans-human Inequity (2038) 367
C H A P T E R 7 9 370
A Toolkit for the Concerned Citizen 370
C H A P T E R 8 0 374
A Documentation Guide for Affected Individuals 374
C H A P T E R 8 1 378
Building Local Networks of Solidarity and Mutual Aid 378
C H A P T E R 8 2 383
The Choice Before Us 383
C H A P T E R 8 3 385
A Call to Action — What You Can Do 385

PART I

Foundations — Understanding Neurotechnology and Neuro-rights

The human mind has long been considered the last private place—a fortress inaccessible from the outside, where thoughts form freely and consciousness remains sovereign. That assumption is now being challenged by a class of technologies that can, for the first time in history, establish a direct connection between the human brain and external devices. These technologies, collectively known as neurotechnologies, are advancing rapidly, and their implications for human rights are profound.

Neurotechnology encompasses an array of devices and systems that interact with the central nervous system through electrical, magnetic, optogenetic, and other means. Some are designed to understand the brain's functioning; others aim to directly intervene in mental processes to restore lost functions or enhance cognitive capabilities. The particularity of neurotechnologies is that they provide a gateway to interfere with mental and cognitive functions in ways that were previously impossible. They can expose cognitive processes, directly alter a person's mental states and thoughts, bypass an individual's conscious control or awareness, and enable non-consensual external access to thoughts, emotions, and mental states. All of this is fueled by "neurodata"—the information gathered from the brain and nervous system, including inferences extracted from it.

As the United Nations Human Rights Council Advisory Committee observed in its landmark 2024 report (A/HRC/57/61), neurotechnologies are "unique and socially disruptive" precisely because of this capacity to penetrate the inner realm of thought. The report notes that for a long time, the human mind and conscience were considered "the container of the inner self, a fortress, inaccessible from the outside." Penetrating the brain, gaining access to inner mental processes, or directly altering them, represented a frontier not to be crossed. That frontier is now being crossed daily.

Neuro-rights emerge as a necessary response to this new reality. They are not abstract philosophical concepts but practical protections grounded in existing human rights law—freedom of thought, privacy, personal integrity, health, and freedom from torture—that must be reinterpreted and reinforced in the age of brain-reading and brain-writing technologies. The five core neuro-rights—mental privacy, personal identity, free will, fair access to cognitive enhancement, and protection from algorithmic bias—represent the minimum framework needed to ensure that neurotechnology serves humanity rather than subjugates it.

The urgency of this moment cannot be overstated. The neurotechnology industry has grown from $331 million in private investment to $7.3 billion in just ten years. Large-scale brain initiatives have been launched by multiple countries since 2013, and the private sector increasingly drives development toward direct-to-consumer devices. Yet this explosive growth is occurring in a highly deregulated environment. As the UN report warns, "existing loopholes in regulations, lack of technical expertise and capacity and the absence of adequate oversight bodies are factors that will certainly be exploited by large companies seeking profits." Without guardrails, the industry will continue growing unfettered, prioritizing profitability and convenience over ethical and human rights considerations.

The convergence of artificial intelligence with neurotechnology compounds these risks. AI enables the processing and analysis of vast datasets of neurodata, making it possible to decode mental states, predict behavior, and infer intimate characteristics about individuals—their sexual orientation, personality traits, cognitive performance, intentions, beliefs, and emotions. While true "mind-reading" is not yet achievable, current neuroimaging techniques can already make inferences across a broad spectrum of mental states. The potential for misuse is enormous.

This guide is written for citizens who may soon encounter neurotechnology in their workplaces, schools, and homes; for legislators tasked with crafting laws to protect their constituents; and for the global community that must decide whether to establish international frameworks for neuro-rights before the technology outpaces governance. The decisions made in the next few years will determine whether neurotechnology becomes a tool for liberation or oppression—whether it heals or harms, empowers or enslaves.

Understanding what neurotechnology is, what neuro-rights are, and why they matter now is the essential foundation for everything that follows.

CHAPTER 1

What Is Neurotechnology? A Layperson's Guide

Neurotechnology is not science fiction. It is already here, in doctors' offices, research laboratories, and increasingly in consumer products sold over the counter. But what exactly does the term mean?

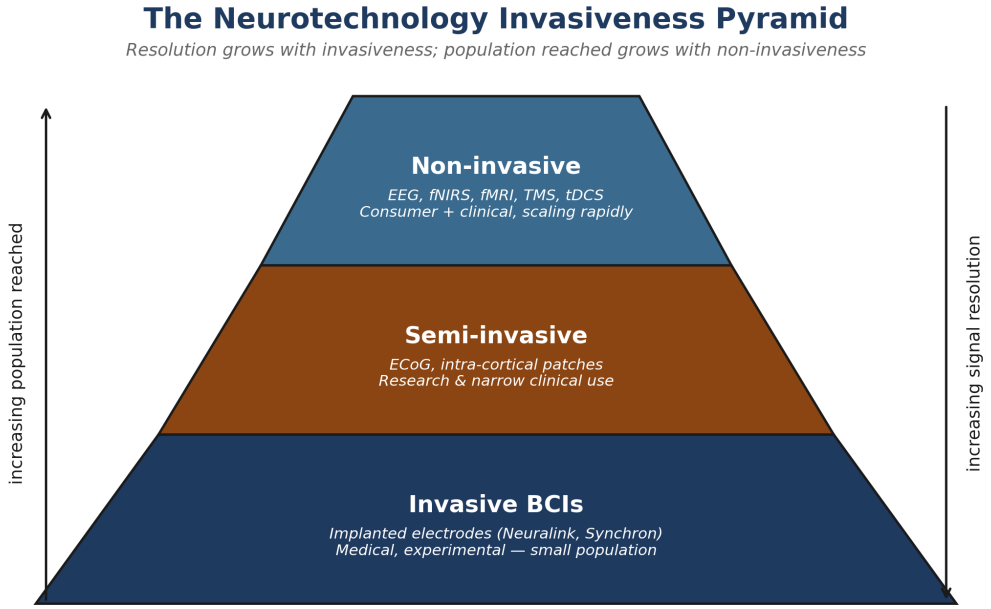


Figure 1.1. *The neurotechnology invasiveness pyramid — signal resolution rises with invasiveness, but the population reached rises with non-invasiveness. Most consumer-facing concern lies in the broad top tier.*

Defining Neurotechnology

The United Nations Human Rights Council Advisory Committee, in its landmark 2024 report (A/HRC/57/61), offers a clear definition: neurotechnology encompasses "an array of devices and systems that interact with the central nervous system through electrical, magnetic, optogenetic and other means." Some of these devices are designed primarily to understand how the brain functions, while others "consist of methods aimed at directly intervening in mental processes with the purpose of restoring lost functions and enhancing cognitive capabilities."

For centuries, the human mind was considered a fortress—inaccessible from the outside. Neurotechnology changes this fundamental reality. As the UN report explains, these technologies "allow a direct connection to be established between the human brain and external devices and thus provide a gateway to interfere with mental and cognitive functions."

What Makes Neurotechnology Unique

Neurotechnology is not just another form of digital technology. The UN report identifies six characteristics that make it "unique and socially disruptive":

1. It can expose cognitive processes that were previously private
2. It can directly alter a person's mental processes and thoughts
3. It can bypass an individual's conscious control or awareness
4. It enables non-consensual external access to thoughts, emotions, and mental states
5. It relies on "neurodata"—brain data needed for its own functioning, calibration, and optimization
6. It collects, analyzes, and processes large personal datasets of a highly sensitive nature

Two Basic Types: Reading and Writing

Within the medical realm, neurotechnologies fall into two basic categories.

Neuroimaging devices observe the structure or functioning of the brain, mapping or measuring its activity. These techniques detect neural signals and allow interpretation of mental states and behaviors—essentially, they *read* the brain. They cannot generally modify the processes they observe.

Neuromodulation techniques and devices do the opposite: they alter and modulate the functioning of the brain by inhibiting or stimulating targeted parts of the brain. These devices *write into* the brain.

The Brain-Computer Interface

An important emerging category is the brain-computer interface (BCI). The UN report defines BCIs as "devices and methods that establish a direct connection and/or communication between the brain and external devices (i.e. phone or computer), allowing the possibility of exchanging information between the two."

A BCI works by collecting the brain's neural activity using recording devices, decoding these activities to extract information about the individual, and using this information to output corresponding commands that control external devices. This creates what engineers call an "interactive closed loop"—the device reads brain signals, processes them, and responds in real time.

Invasive vs. Non-Invasive Technologies

The physical invasiveness of a neurotechnology determines both its power and its risks.

Invasive applications require surgical implantation of electrodes or implants into the brain. They are more precise and powerful but entail significant risks—surgery, infection, and long-term complications. Deep brain stimulators for treating Parkinson's disease have been implanted worldwide for decades. More recently, companies have begun large-scale trials of chip technology that can be implanted into the brain through blood vessels, while others advertise "cosmetically invisible" implants that allow users to control computers from any location.

Non-invasive neurotechnologies do not penetrate skin, skull, or brain tissue. They are used externally through various supports: helmets, glasses, wristbands, electrode patches, tattoos, and earphones. The new generation of wireless earbuds includes biosensing electrodes that can measure brain activity. These devices are far less risky but also less precise.

Consumer Devices on the Market

Non-invasive neurotechnologies are being commercialized for a wide range of non-medical applications: improving cognitive and sporting performance, aiding concentration, promoting relaxation, and enhancing sleep. Users can wirelessly track their levels of concentration and check results on their phones.

However, the UN report issues a critical warning: "The effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge. Moreover, research on potential side-effects, long-term impacts and their possible irreversible character is lacking."

A Rapidly Growing, Deregulated Industry

The scale of neurotechnology development is staggering. The UN report notes that large-scale "brain initiatives" have been launched by the United States, Australia, China, the European Union, Israel, Japan, and the Republic of Korea since 2013. Private investment has grown from $331 million to $7.3 billion in just ten years.

Yet this growth is happening largely without regulation. While medical neurotechnologies are governed by established frameworks, "the consumer sector remains highly deregulated and is characterized by rapid expansion." The UN report warns that products with inadequate safety measures and unclear human rights risks "may become pervasive throughout daily life despite the fact that, in most countries, applicable regulations are unclear, weak or non-existent."

Glossary of Key Terms

Brain-computer interface (BCI): A device that establishes direct communication between the brain and an external device such as a phone or computer.

Closed-loop system: A neurotechnology that can both read and write neural signals, adapting its function autonomously based on feedback.

Invasive neurotechnology: A device that requires surgical implantation of electrodes or implants into the brain.

Neurodata: Information gathered from the brain and/or nervous system, including inferences extracted from it.

Neuroimaging: Techniques used to observe the structure or functioning of the brain and map or measure its activity.

Neuromodulation: Techniques that alter brain function through inhibition or stimulation of targeted brain regions.

Neurotechnology: Devices and systems that interact with the central nervous system through electrical, magnetic, optogenetic, or other means.

Non-invasive neurotechnology: A device that does not penetrate skin, skull, or brain tissue, used externally through wearable supports.

Understanding these fundamentals is essential before examining the specific rights at stake, the technologies already in use, and the urgent need for legal frameworks to protect what may be humanity's last private place—the human mind.

CHAPTER 2

The Spectrum of Neurotechnologies — From EEG Headbands to Brain Implants

Neurotechnology is not a single invention but a rapidly expanding family of devices and systems that interact with the central nervous system through electrical, magnetic, optogenetic, and other means. As the 2024 United Nations Human Rights Council report explains, these technologies "generally enable the exposition of cognitive processes" and "enable the direct alteration of a person's mental processes and thoughts." Understanding the spectrum of existing technologies is essential for grasping both the promise and the peril they represent.

The Two Fundamental Functions: Reading and Writing

At their most basic level, neurotechnologies perform one of two functions—or both. The UN report draws a clear distinction between **neuroimaging**, which involves "devices and techniques used to observe the structure or functioning of the brain and map or measure its activity," and **neuromodulation**, which alters brain function "through the inhibition or stimulation of targeted parts of the brain." Neuroimaging is the capacity to *read* neural signals; neuromodulation is the capacity to *write* into the brain.

When a device can do both, it becomes a **closed-loop system**—one that can "adapt itself autonomously" based on feedback from the brain it is simultaneously reading and modifying.

Non-Invasive Technologies: The Wearable Revolution

The most widely available neurotechnologies are non-invasive—they "do not penetrate skin, skull or brain tissue" and are used externally through various supports: "helmets, glasses, wristbands, electrode patches, tattoos and earphones." The UN report notes that "the most popular wearables are used to improve cognitive and sporting performance or to aid concentration, relaxation or sleep." Users can wirelessly track their levels of concentration and check results on their phones.

Electroencephalography (EEG) headsets are the most common non-invasive devices. They use sensors placed on the scalp to detect electrical activity in the brain. Consumer-grade EEG devices are marketed for meditation, focus training, and even gaming. The UN report cautions, however, that "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge" and that "research on potential side-effects, long-term impacts and their possible irreversible character is lacking."

A newer frontier is the integration of biosensing electrodes into everyday audio devices. "The sensor systems of the new generation of wireless earbuds include biosensing electrodes, enabling the measurement of various biosignals, including brain activity." This means consumers may soon be wearing neurotechnology without even knowing it.

Functional magnetic resonance imaging (fMRI) is a non-invasive neuroimaging technique that measures brain activity by detecting changes in blood flow. While primarily a medical and research tool, fMRI is increasingly explored for applications beyond diagnosis. The UN report notes that neuroimaging "can be employed to make inferences about individuals on a large spectrum of mental states, including memories, semantic knowledge, emotions, dreams, inner speech and intentions."

Transcranial magnetic stimulation (TMS) uses magnetic fields to stimulate nerve cells in the brain. It is non-invasive and has been approved for treating depression. TMS represents neuromodulation without surgery—a way to write into the brain from outside the skull.

Invasive Technologies: Surgical Precision and Greater Risk

Invasive neurotechnologies "require the surgical implantation of electrodes or implants" and are "more precise and powerful but entail more risks." The most established invasive technology is **deep brain stimulation (DBS)** , which "has been in use for decades and is being implanted worldwide for the treatment of neurological conditions" such as Parkinson's disease, epilepsy, and chronic pain.

The most ambitious invasive technologies are **brain-computer interfaces (BCIs)** , which "establish a direct connection and/or communication between the brain and external devices (i.e. phone or computer), allowing the possibility of exchanging information between the two." The UN report describes the process: "It involves collecting the brain's neural activity using recording devices, decoding these activities to extract information about the individual and using this information to output corresponding commands to control external devices, thereby creating an interactive closed loop."

Recent developments have pushed BCIs beyond the purely medical realm. The UN report notes that "a company that has recently developed a secure interface for communication 'with the power of thought' is conducting large-scale trials of this technology, which can be implanted into the brain through blood vessels." Other companies "already advertise 'cosmetically' invisible implants that may allow users to control computers or mobile devices from any location" and "claim to be working on the 'cognitive enhancement, pleasure or leisure of healthy people.'"

The BCI as Gateway

The brain-computer interface represents a qualitative leap in neurotechnology. Unlike earlier devices that simply monitored or stimulated brain activity, BCIs create a two-way communication channel. The UN report emphasizes that "the particularity of neurotechnologies is that they allow a direct connection to be established between the human brain and external devices and thus provide a gateway to interfere with mental and cognitive functions."

This gateway is what makes BCIs so promising for restoring function in paralysis or locked-in syndrome—and so dangerous if misused. With the support of artificial intelligence and

machine learning, “the opportunities and uses of neurotechnologies are multiplying.” AI enables the processing and analysis of large sets of neurodata, making it possible to decode increasingly complex mental states.

Emerging Technologies and the Convergence with AI

The UN report notes that “modern neurotechnologies are gaining in complexity as a result of their convergence with other fields, such as engineering and computational sciences.” This convergence is accelerating development in several directions:

Optogenetics uses light to control neurons that have been genetically modified to respond to light. This allows unprecedented precision in targeting specific neural circuits.

Nanotechnologies are being developed that could interact with the nervous system at the cellular level, potentially enabling non-surgical access to deep brain structures.

Closed-loop systems that combine reading and writing capabilities are becoming more sophisticated, raising the prospect of devices that can automatically adjust a person’s mental state without their conscious awareness.

The Deregulated Consumer Frontier

The UN report draws a stark contrast between medical and consumer neurotechnologies: “In contrast to the medical development of neurotechnologies, which are generally regulated by established frameworks and processes, the consumer sector remains highly deregulated and is characterized by rapid expansion.”

The report warns that “the possibility that, in the coming years, those products with inadequate safety measures and unclear or underestimated human rights risks may be widely commercialized is real.” The risk is that “without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations.”

From EEG headbands marketed for better meditation to implants promising thought-controlled computing, the spectrum of neurotechnologies spans devices that are already in millions of homes and systems that remain experimental. What unites them is their capacity to reach into the last private place—and the urgent need for frameworks that ensure this reach serves human dignity rather than undermining it.

CHAPTER 3

How the Brain Actually Works — A Neuroscience Primer

Before we can meaningfully debate the rights of citizens against neurotechnology, we must understand, in plain language, what the technology is acting upon. The brain is the most complex object known in the universe, and the language used to describe it can obscure rather than illuminate. This chapter offers a layperson's primer on the structure and function of the human nervous system, the signals that neurotechnology measures and modulates, and the levels at which those technologies operate. It is not exhaustive. It is sufficient to follow the legal, ethical, and political arguments that fill the rest of this book.

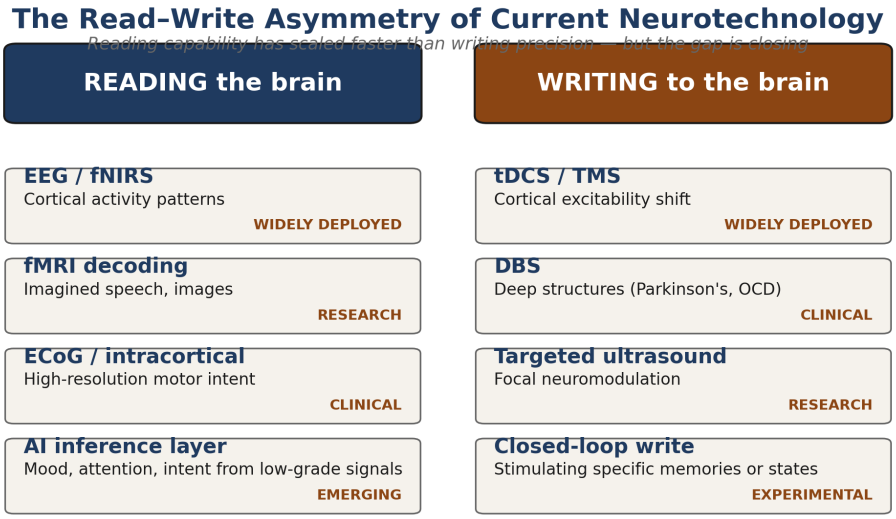


Figure 3.1. *The read–write asymmetry of current neurotechnology. Reading capability has scaled faster than writing precision, but the gap is narrowing as closed-loop systems mature.*

The Architecture of the Nervous System

The human nervous system is divided into two large compartments. The central nervous system comprises the brain and the spinal cord. The peripheral nervous system comprises the network of nerves that radiates outward from this central core, carrying signals to and from muscles, organs, sensory receptors, and skin. Neurotechnology, in its strict sense, targets the central nervous system: implants in the brain, electrodes on the scalp, magnetic coils held over the skull, optical probes that read changes in blood oxygenation in cortex. But because the peripheral nervous system is the gateway through which the central one experiences the world, peripheral neurotechnologies — vagus nerve stimulators, retinal implants, cochlear devices — are part of the same regulatory conversation.

The brain itself is roughly the size of two fists pressed together and weighs about 1.4 kilograms. Yet within this small volume reside approximately 86 billion neurons, each connected to thousands of others, generating a network whose total number of connections — roughly 100 trillion — exceeds the number of stars in the Milky Way by a factor of a thousand. It is this connective density, more than the number of cells, that gives the brain its computational power and its mystery.

The Major Regions and What They Do

It is useful, though always a simplification, to describe the brain as a collection of regions specialized for different functions. The cerebral cortex — the outer wrinkled sheet of tissue — is the seat of conscious thought, language, planning, voluntary movement, and most of the cognitive functions that distinguish humans from other species. Underneath the cortex lie the basal ganglia, involved in habit, motor control, and reward; the thalamus, which acts as a relay station for nearly all sensory information; the hypothalamus, which regulates the body's internal state and links the nervous system to the hormonal one; and the brainstem, which controls heart rate, breathing, and the basic functions that keep us alive.

Within the cortex, four lobes carry the bulk of the functional load. The frontal lobe handles executive functions — planning, decision-making, working memory, and the suppression of impulses. The parietal lobe processes spatial information and bodily sensation. The temporal lobe processes auditory information, language, and memory. The occipital lobe processes vision. Each lobe contains many specialized sub-areas, and any complex behavior — reading a sentence, recognizing a face, formulating a moral judgment — recruits multiple regions in coordination.

This functional map matters for neurotechnology because the technology's effectiveness depends on whether it can reach the right tissue. A deep-brain stimulator implanted in the subthalamic nucleus can reduce the tremors of Parkinson's disease because the target sits deep in the brain and is reachable only by surgical implantation. An electrode on the scalp picks up signals from the surface of the cortex but cannot easily resolve activity in deep structures. A focused ultrasound device can reach deeper tissue without surgery but at a lower spatial resolution. Different technologies see and touch different parts of the brain, and the ethical and legal stakes vary accordingly.

Neurons and Their Signals

The basic computational unit of the nervous system is the neuron — a single cell that receives, integrates, and transmits information. A neuron has three principal parts: dendrites, which receive input from other neurons; the cell body, which integrates that input; and an axon, which transmits an output signal to other cells. Communication between neurons is both electrical and chemical. Within a single neuron, signals travel as electrical impulses called action potentials. Between neurons, signals are passed across small gaps called synapses by chemical messengers known as neurotransmitters.

When a neuron is active, it generates a brief electrical event — a spike — lasting about a millisecond. The simultaneous activity of millions of neurons produces larger rhythmic

oscillations that can be detected from outside the skull. These rhythms, organized by frequency band, are what an electroencephalogram (EEG) measures. Slow oscillations in the delta band (1–4 hertz) dominate during deep sleep. Theta rhythms (4–8 Hz) accompany memory encoding. Alpha rhythms (8–13 Hz) emerge when the eyes close and the brain enters a relaxed waking state. Beta (13–30 Hz) and gamma (30–100 Hz) accompany attention, cognitive effort, and conscious perception. These bands are the alphabet from which all electrical brain signatures are constructed, and consumer EEG devices are essentially low-resolution alphabet readers.

The Vascular Signal

Neurons consume energy when they fire. To meet this demand, local blood vessels dilate within seconds, delivering oxygenated blood to active regions. This haemodynamic response is the basis of functional magnetic resonance imaging (fMRI). An fMRI scanner detects the magnetic signature of oxygenated versus deoxygenated blood, allowing researchers to map activity across the entire brain with millimetre-scale spatial precision. The trade-off is temporal: the haemodynamic response lags actual neural activity by several seconds. Where EEG reads the brain's electrical signature in milliseconds but cannot say precisely where the signal comes from, fMRI reads the location with high precision but smears the timing.

Functional near-infrared spectroscopy (fNIRS), a much more portable technology, exploits the same vascular signal using light rather than magnetic fields. Light in the near-infrared range passes through skin and skull and is differentially absorbed by oxygenated and deoxygenated haemoglobin. By placing emitters and detectors on the scalp, an fNIRS device can map cortical blood-flow changes without confining the subject to a giant magnet. fNIRS is now sufficiently cheap and small to be deployed in classrooms, workplaces, and consumer products, raising privacy concerns that the technology's MRI ancestor never did.

Reading and Writing

Neurotechnologies fall into two functional categories: those that read the brain and those that write to it. Reading technologies passively observe the signals the brain produces. Writing technologies actively introduce energy or chemistry into the brain to alter its activity. The distinction is critical because the ethical stakes differ. Passive observation is a privacy concern. Active modulation is a personal-integrity concern. The most powerful systems do both — they read the brain, decide how to modulate it, then write back into it, all in a closed loop.

Reading technologies include EEG, magnetoencephalography (MEG, which detects the tiny magnetic fields generated by neural currents), fMRI, fNIRS, electrocorticography (ECoG, in which electrodes are placed directly on the cortical surface during neurosurgery), and microelectrode arrays implanted within brain tissue. Each measures different aspects of brain activity at different scales. EEG measures synchronized populations of cortical neurons at scalp level. MEG resolves source locations more precisely. ECoG captures local field potentials beneath the dura. Microelectrode arrays record individual neuron spikes.

Writing technologies include transcranial magnetic stimulation (TMS), which uses brief magnetic pulses to induce currents in cortical tissue; transcranial direct-current stimulation (tDCS) and transcranial alternating-current stimulation (tACS), which pass weak currents between scalp electrodes; transcranial focused ultrasound (tFUS), which focuses acoustic energy into deep tissue; deep brain stimulation (DBS), which delivers electrical pulses through implanted electrodes; and various drug-delivery and optogenetic techniques used in research. Each technology operates at a different depth, resolution, and intensity. Some are reversible. Others are not.

Why Resolution Matters

When debating neuro-rights, the resolution of a technology is often more important than the technology's name. A high-resolution reading device can extract more information from neural activity — specific words being heard, the identity of a remembered face, the emotional valence of an experience — and therefore poses a greater privacy threat. A high-resolution writing device can induce more specific changes — a particular mood, a craving, a forced motor command — and therefore poses a greater integrity threat. The same underlying technology can be benign at low resolution and alarming at high resolution. EEG that distinguishes only sleep from wakefulness is unremarkable. EEG that decodes inner speech is something else entirely.

Resolution improves on three axes: spatial (how precisely the device localizes activity in the brain), temporal (how precisely it tracks the timing of activity), and inferential (what an algorithm can deduce from the recorded signals, given a sufficient training set). It is the third axis — the algorithmic inference — that has grown most dramatically in the last decade and that the law most poorly addresses. The same EEG signal that revealed only crude arousal in the 1980s can today be decoded into a probability that the subject is depressed, is concentrating, has recognized a familiar image, or is recalling a specific category of memory. The hardware did not change; the decoders did.

Three Axes of Neural Resolution

Spatial resolution. How small a region of brain can the device distinguish? EEG: centimetres. fMRI: millimetres. Microelectrode arrays: micrometres.

Temporal resolution. How quickly can the device track changes? EEG and MEG: milliseconds. fMRI and fNIRS: seconds.

Inferential resolution. What can be decoded from the signal? Driven by algorithms, training data, and computing power — not by hardware. This axis is the one that has changed most dramatically, and is the one current law least restrains.

The Blood-Brain Barrier

One feature of brain biology shapes nearly every debate about neurotechnology: the blood-brain barrier. This biological membrane, formed by tightly joined cells lining brain capillaries, prevents most substances in the blood from entering the brain. It is the reason that many drugs which would be effective in a Petri dish do nothing when swallowed, and the reason that pharmacological control of brain function is so much less precise than its inventors hoped. The blood-brain barrier is also the reason that many neurotechnologies seek non-pharmacological pathways into the brain — electrical, magnetic, optical, acoustic. Each of these energy modalities has its own absorption characteristics through the scalp and skull, and each interacts differently with biological tissue. None is as fully understood, in terms of long-term safety, as the public is generally told.

From Biology to Rights

This is the biological substrate that neuro-rights propose to protect. The brain is not just an organ; it is the physical seat of consciousness, identity, and agency. Its signals are not simply data; they are the trace of mind. When a technology reads or writes to this substrate, it does not merely process information — it touches a person at the deepest accessible level of their being. The legal and ethical structures the rest of this book describes are attempts to draw lines around that touch — lines that say, here are the conditions under which access is permitted, and beyond them, access is unjust.

Understanding the biology is the first step in understanding why those lines matter. A debate about neural privacy that has no model of the underlying biology will inevitably misallocate concern. It will worry about technologies that pose little real threat and ignore those that pose great ones. It will set rules that are easy for sophisticated actors to circumvent. It will reassure citizens with protections that, viewed against the technology's true capabilities, are essentially decorative. To avoid that fate, citizens, legislators, and judges all need a working grasp of what the brain is doing and how the technology engages with it. The chapters that follow build on the foundation laid here.

CHAPTER 4

Reading and Writing the Brain — Decoding and Encoding

Of all the metaphors used to describe neurotechnology, none is more useful than the pair of verbs: reading and writing. The brain is, at one level, a generator of electrical and chemical patterns. To read the brain is to detect those patterns and infer what they mean. To write to the brain is to inject patterns of one's own and see how they alter what the brain does next. Together, reading and writing form the technical foundation of every claim and counter-claim that fills the legal debates of this book. This chapter examines, in some depth, what reading and writing actually involve — how each is done, what each can and cannot accomplish today, and what each is likely to be able to do tomorrow.

The Reading Pipeline

Reading the brain is rarely a single step. It is a pipeline. Raw neural activity must first be acquired by a sensor. The sensor produces a noisy, redundant electrical record that must be cleaned of artifacts — muscle movement, eye blinks, environmental electromagnetic interference. The cleaned record is then transformed into features that a decoding algorithm can use. Those features are passed to a model that has been trained, on previously labelled data, to associate particular feature patterns with particular mental states, intentions, or perceptual contents. The model outputs a probability distribution over possible interpretations. A confidence threshold determines which interpretation is reported. The entire pipeline runs in milliseconds.

Each step in the pipeline introduces opportunities for both privacy violations and errors. A noisy sensor records signals from neighbouring brains as well as from the target. A poorly chosen feature extractor discards relevant information. A model trained on a non-representative population overgeneralizes. A confidence threshold set too low produces hallucinated interpretations; set too high, it produces silences mistaken for absences. Most regulatory debate focuses on the sensor and the output. The middle of the pipeline — the feature extractors and the trained models — is where most of the real ethical action takes place, and is the part the law most often ignores.

What Can Currently Be Decoded

Public discussion of neural decoding tends to oscillate between two errors. Some commentators claim that mind-reading is here — that scientists can already extract the contents of consciousness on demand. Others insist that nothing of the sort is possible — that brain signals are too noisy and too individual to permit meaningful interpretation. Both positions are wrong. The reality is more nuanced and more troubling than either.

Today's decoders can reliably distinguish, from neural signals alone, a range of categories. They can determine whether a subject is awake or asleep, attending or distracted, viewing a face or a house, performing arithmetic or planning a movement. They can decode the

trajectory a hand is intended to make from microelectrode arrays in motor cortex, allowing paralysed patients to control robotic limbs and computer cursors. They can identify which of several pre-trained categories a subject is thinking about. With training-set sizes that grow each year, they can recognize specific words from inner speech in cooperating subjects, recover blurry but recognizable images from visual cortex activity, and predict whether a subject will remember a piece of information later.

What decoders cannot do, today, is read arbitrary thoughts from an uncooperative subject. The training process required is non-trivial: a model trained on subject A's neural signals usually performs poorly on subject B. Without subject cooperation and a calibration session, current decoding accuracy degrades sharply. Without informed cooperation — the subject knowing what is being decoded and adjusting their thinking accordingly — accuracy degrades further. Mind-reading, in the pure sense of the science-fiction trope, is not yet possible.

This limitation, however, is rapidly eroding. Larger training datasets, transfer-learning techniques that adapt one subject's model to another, and the use of foundation models that share representations across many individuals all push toward subject-independent decoding. The most plausible projection is that within the next decade, decoders will operate with usable accuracy on naive subjects under controlled conditions, and within two decades, on naive subjects under uncontrolled conditions. The law has perhaps ten years to catch up with capabilities that will outrun it.

Decoding Task	Current Status	5-Year Projection
Sleep/wake state	Trivial; consumer EEG devices do it	Refined to sub-state granularity
Attention and arousal	Reliable in trained subjects	Reliable without calibration
Movement intention	Well established for trained users	Real-time, naturalistic control
Visual category	Reliable for trained categories	Open-set recognition
Inner speech (specific words)	Demonstrated in cooperating subjects	Continuous decoding in cooperators
Emotional valence	Coarse categories with training	Fine-grained, real-time
Specific memories	Limited, with prior labelling	Open-set retrieval in cooperators
Identity, beliefs, intentions	Statistical inference only	Direct decoding becomes plausible

The Inference Layer

Most legal debate about neural data conceives of it as a kind of biometric — a string of measurements taken from the body. But the more important risks lie in what is inferred from those measurements, often in combination with other data. A neural signature alone may reveal little. The same signature combined with the time of day, the application being used, the content the subject is viewing, and the metadata of the device, can permit inferences far beyond what the raw signal contains. The legal definition of neural data, in jurisdictions that have one, must include these inferences — or it protects nothing. As one drafting note appended to the EU petition produced through TI One Voice puts it: the single highest-value legislative fix is defining neural data to include the algorithmic inferences drawn from it, not only the raw signal.

The Writing Pipeline

Writing to the brain is, in its current state, less precise than reading but more consequential. Reading errors are misinterpretations. Writing errors are interventions in a living organ. A misread thought is a privacy violation, perhaps a slander. A misexecuted stimulation is, at minimum, a physical event with unpredictable consequences.

Writing technologies fall into two broad classes. Indirect writing alters neural function by changing the chemical or physiological context in which neurons operate — pharmacological agents, hormones, dietary interventions. Direct writing introduces energy into neural tissue itself — electrical, magnetic, optical, or acoustic — to drive neurons to fire or not to fire. Indirect writing is older and better understood. Direct writing is newer, more rapidly evolving, and more concerning from a rights perspective because it can be applied without ingestion and, in some configurations, without the subject's knowledge.

Transcranial magnetic stimulation produces brief, intense magnetic pulses that induce currents in cortical tissue. A single pulse to motor cortex can produce an involuntary muscle twitch in the corresponding limb. Sequences of pulses, delivered with carefully tuned timing, can modify the excitability of a cortical region for minutes to hours after the stimulation ends. TMS is approved for the treatment of depression that has not responded to other therapies and is used experimentally in dozens of other conditions. It is also, in principle, applicable to anyone, with or without consent, who can be brought within range of a magnetic coil.

Transcranial focused ultrasound, the newest entrant in the field, can deliver mechanical energy with millimetre precision to targets deep in the brain, without crossing the skull surgically. It is a powerful research tool. It is also, like TMS, deployable at a distance from the subject. The combination of focused reach and non-invasive application is one of the developments that has most accelerated the regulatory debate this book chronicles.

Closed-Loop Systems

The most consequential class of neurotechnology is the closed loop: a system that reads the brain, decides how to respond, and writes back, in a continuous cycle, often with no human in the decision step. Closed loops have legitimate medical uses. A closed-loop deep brain stimulator can detect the neural signature of an oncoming seizure and deliver a counter-

stimulus before the seizure becomes clinical. A closed-loop neuroprosthesis can sense an intended movement and execute it via a robotic arm. These applications restore function lost to disease or injury, and they are among the most defensible uses of neurotechnology.

But closed loops can also be designed to act on cognitive states rather than motor ones, and to optimize for goals other than the patient's welfare. A device that detects the neural signature of consumer indecision and triggers a marketing prompt is technically feasible. A device that detects the signature of dissent and triggers an aversive stimulus is, in laboratory conditions, demonstrable. A device that monitors a worker for the signature of cognitive fatigue and triggers a productivity intervention is already plausible. The line between medical assistance and behavioural manipulation is, at the technical level, a single line of code.

The Asymmetry Between Read and Write

Reading and writing differ in one further critical respect: their detectability. A reading device emits, at most, a brief electromagnetic signature when it transmits collected data. A passive recording device may emit nothing at all. A writing device, by contrast, must deliver energy into the subject, and that energy is, in principle, measurable. A pulsed magnetic field, an ultrasonic beam, a directed millimetre-wave emission — each leaves a physical signature that appropriately equipped detection can find. This asymmetry has practical implications for the legal and forensic frameworks discussed in later chapters. Privacy violations through reading are extraordinarily difficult to prove because the act of reading leaves few traces. Integrity violations through writing are difficult to prove in practice, because detection requires specialized equipment, but they are not in principle undetectable. The forensic case for regulating writing is, on the evidence alone, considerably stronger.

The Limits of the Metaphor

Reading and writing are metaphors. They oversimplify. The brain is not a text; its activity does not consist of discrete words waiting to be retrieved. The signals neurotechnology measures are statistical regularities across populations of cells, summarized by algorithms that introduce their own structure. What the algorithm reports as a thought is not a thought; it is a probability that a thought of a certain type is present, conditioned on a model trained on data from other subjects. To take the algorithm's output at face value is to confuse the map for the territory. To dismiss it as unreliable, however, is to ignore that the map is increasingly accurate, that maps are sufficient for many purposes a citizen would prefer to refuse, and that adversaries operating without scruples will use the maps regardless of philosophical reservations. The job of legal protection is not to wait until the map is perfect; it is to set the terms of map-making while the cartographers can still be regulated.

CHAPTER 5

What Are Neuro-rights? The Five Core Rights Explained

If neurotechnology can reach into the mind, what rights protect the person whose mind it is? This question has driven a global effort to define a new category of human rights—neuro-rights—specifically designed to address the unique capabilities of technologies that interact with the brain.

The concept of neuro-rights emerged from the recognition that existing human rights frameworks, while foundational, were not designed to address technologies that can directly access, decode, or alter mental processes. As the 2024 UN Human Rights Council report (A/HRC/57/61) states, neurotechnologies are "unique and socially disruptive" because they "enable the exposition of cognitive processes," "enable the direct alteration of a person's mental processes and thoughts," and "bypass the individual's conscious control or awareness."

Building on this understanding, advocates and experts have articulated five core neuro-rights. These are not replacements for existing rights but rather specific elaborations of how fundamental human rights must apply in an age of brain-reading and mind-modifying technology.

1. Mental Privacy: The Right to Keep Your Thoughts to Yourself

Mental privacy is the right to control access to your own neural data—the information generated by your brain's activity. The UN report emphasizes that "neurodata provide highly sensitive information and allow aspects of individuals' identity and personality to be revealed, such as sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions."

This right addresses the fundamental concern that, without protection, the contents of a person's mind could be accessed, recorded, analyzed, or sold without their knowledge or meaningful consent. The UN report notes that "consumer neurotechnology companies are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future."

Mental privacy is closely linked to the right to privacy under Article 17 of the International Covenant on Civil and Political Rights, but neurotechnology creates new dimensions of intrusion. As the UN report observes, "Neurotechnologies introduce new threats to the right to privacy, increasing the level of possible interference and impact."

2. Personal Identity: The Right to Be Yourself

Personal identity as a neuro-right means that external technologies should not be able to alter your fundamental sense of self, your personality, or your core beliefs without your

consent. This right is grounded in the principle of human dignity, which the UN report describes as "the essence of what it means to be human."

The UN report notes that neurotechnologies can "affect the constitutive elements of the person: the individual's personhood, mental capacities and personality." Patients who have undergone deep brain stimulation have reported "feeling a changed sense of agency and identity." The report warns that neurotechnologies can "result in actions causing mental harm" and provoke "alterations in an individual's personality, psychological balance or sense of self-identity."

This right protects against what the UN report calls "the objectification and instrumentalization of individuals"—treating a person's mind as something to be modified for external purposes rather than respected as the core of their being.

3. Free Will: The Right to Make Your Own Decisions

Free will as a neuro-right protects the ability of individuals to make decisions free from covert or coercive neural manipulation. The UN report is explicit: "Everyone has the right to develop thoughts autonomously, free from 'impermissible' external influence."

This right is grounded in freedom of thought under Article 18 of the International Covenant on Civil and Political Rights, which the UN report describes as "a far-reaching and profound right that encompasses all matters." The report emphasizes that freedom of thought "does not allow for any interference whatsoever and must be protected unconditionally."

The UN report identifies specific threats to free will from neurotechnology, noting that "neuromarketing strategies have already demonstrated, they can be successfully used to condition the forming of opinions, as well as influencing an individual's decision-making processes." More concerning, "neurotechnologies allow the level of conscious reasoning to be bypassed, leaving individuals without psychological protection from having their mind involuntarily read."

4. Fair Access to Cognitive Enhancement: The Right to Equal Opportunity

As neurotechnologies develop the capacity to enhance cognitive function—improving memory, attention, or processing speed—questions of equitable access become critical. The UN report recognizes that "when safe, effective, secure and human-rights compliant neurotechnology products do exist, access becomes a key element of the right to health."

This right addresses the concern that cognitive enhancement technologies could create a "neuronal divide" between those who can afford enhancement and those who cannot. The UN report warns that "persons with disabilities, in particular, must be granted access to affordable assistive technologies without discrimination," and that "universal access to advanced technology depends, however, on needs assessments and budgetary decisions."

The right to fair access is not a guarantee that everyone receives the same enhancement, but rather that access is not determined solely by wealth or social status, and that vulnerable

populations are not excluded from beneficial technologies while being disproportionately exposed to risky ones.

5. Protection from Algorithmic Bias: The Right to Fair Treatment

As neurotechnologies increasingly rely on artificial intelligence to interpret neural signals and make decisions, the risk of algorithmic bias becomes acute. The UN report notes that "issues resulting from the convergence with artificial intelligence, such as algorithmic bias and lack of traceability and transparency, will also be exacerbated through their use."

This right protects against discrimination based on neural data. The UN report warns that "neuroprofiles may be used to discriminate among people on the basis of their mental signatures, leading to decisions (i.e. in relation to the workplace or insurance coverage) that may have a negative impact on the interests and rights of the data subject."

The report emphasizes that "persons with disabilities are particularly exposed to medical and scientific procedures, research and experimentation" and that "without meaningful engagement and participation of persons with disabilities in their design, development and implementation, neurotechnologies risk contributing to stigmatization and objectification, reinforcing structural-based discrimination."

The Foundation of All Neuro-rights

Underlying these five rights is the foundational principle of human dignity. The UN report states that neurotechnologies "should first be examined against the foundational principle of human dignity, which encapsulates the essence of what it means to be human." This requires assessing whether neurotechnologies "contribute to the objectification and instrumentalization of individuals" and whether they "hinder individuals' capacity for self-determination."

The five neuro-rights are not abstract concepts—they are practical protections designed to address specific capabilities of neurotechnology that existing human rights frameworks were not designed to handle. As the UN report concludes, "Where areas not sufficiently covered by existing standards do emerge, action is needed to prevent violations, misuses and abuses."

CHAPTER 6

Cognitive Liberty — The Foundational Freedom

Before any of the five neuro-rights can be protected, one deeper principle must be secured: the right to control one's own consciousness. This is cognitive liberty—the foundational freedom upon which all other neuro-rights depend.

What Cognitive Liberty Means

Cognitive liberty is the right of every individual to determine their own mental states, to think their own thoughts, and to make their own decisions about what enters their mind. It is the freedom to alter one's consciousness through whatever means one chooses—meditation, education, substances, or neurotechnology—and equally the freedom to refuse any external manipulation of one's mental processes.

The concept draws from the principle that the *forum internum*—the inner realm of thought—has always been treated as inviolable. Article 18 of the International Covenant on Civil and Political Rights protects freedom of thought as an absolute right, admitting no exception or derogation even in times of war. The reason it could be absolute was practical: 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 inner forum is becoming technically accessible.

Historical Roots

The term "cognitive liberty" was developed within the neuroethics community as a framework for thinking about mental autonomy in the age of psychopharmacology. But its roots run deeper. The right to control one's own mind has been implicit in Western legal traditions since at least the Enlightenment, and in Eastern traditions for millennia. What changed is that what was once a philosophical assumption must now be a legal guarantee.

The Modern Relevance

Cognitive liberty becomes urgent when neurotechnology can do three things:

Read mental states. Devices already on sale can decode attention, mood, and intention from brain signals. Laboratories can reconstruct words a person only imagined speaking and images they only saw in their mind's eye. When thoughts can be read from outside, the freedom to keep them private must be explicit.

Alter mental states. Deep-brain stimulation can quiet Parkinson's tremor and treat depression. Transcranial magnetic stimulation can modulate mood and cognition. Closed-loop systems can detect a mental state and respond with stimulation—crossing from therapy into potential manipulation when the objective function is set by someone other than the user.

Bypass conscious deliberation. The most dangerous capability is not reading or writing to the brain, but doing so without the person's awareness. Covert influence, manipulation that operates below the threshold of conscious perception, and systems that shape behavior without the subject's knowledge all violate cognitive liberty at its core.

The Inalienable Dimension

The critical insight about cognitive liberty is that some aspects of it must be inalienable—they cannot be signed away, regardless of consent. Just as a person cannot validly consent to be enslaved, there should be neural interventions that no terms-of-service checkbox can authorize.

This is not a theoretical concern. 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. The terms of service that users click through grant broad licenses to store, process, and often share the data their brains generate.

Cognitive liberty requires that some protections survive consent. The right to not have one's thoughts covertly influenced, the right to not have one's mental states manipulated without awareness—these cannot be conditional on a user's willingness to accept a company's terms.

The Relationship to Other Neuro-rights

Cognitive liberty is the foundation because it encompasses the others. Mental privacy is cognitive liberty's defensive aspect—the right to keep thoughts private. Personal identity is cognitive liberty's continuity aspect—the right to remain oneself over time. Free will is cognitive liberty's agency aspect—the right to act on one's own decisions. Fair access to cognitive enhancement is cognitive liberty's equality aspect—the right to develop one's mind without being disadvantaged. Protection from algorithmic bias is cognitive liberty's integrity aspect—the right to have one's mental states judged accurately and fairly.

Without cognitive liberty, the other rights become contingent, subject to the permission of whoever controls the technology. With cognitive liberty secured, each of the other rights follows naturally.

What a Real Guarantee Requires

A practical guarantee of cognitive liberty requires four things simultaneously:

A default of inaccessibility. Neural data should be opt-in only, never the price of using a device. The burden of proof must fall on those who seek to read or alter mental states.

Purpose limitation with teeth. Data gathered to help you sleep cannot be repurposed to score you for a loan. The consent you gave was for what the data could reveal then, not what it can reveal now.

Prohibitions that survive consent. Some neural interventions should be forbidden regardless of what any user agrees to—covert influence, manipulation that bypasses conscious deliberation, and any alteration of mental states without the subject's knowledge.

Enforceability. Rights without remedies are decoration. The Chilean Supreme Court's 2023 ruling that ordered a company to delete brain data collected without proper consent showed that enforceable cognitive liberty is possible.

The Philosophical Stakes

Cognitive liberty rests on a claim about human nature: that the capacity to direct one's own mental life is what makes autonomy meaningful. If that capacity can be overridden from outside—if thoughts can be read without permission, if moods can be adjusted without awareness, if decisions can be shaped without consent—then the foundation of democratic citizenship itself is eroded.

A society that protects cognitive liberty is one that acknowledges the inviolability of the inner forum. A society that does not is one that has accepted, whether explicitly or by neglect, that the last private place is open for business.

CHAPTER 7

Why Now? The Urgency of Neuro-rights in the 2020s

The question is not whether neurotechnology will transform society. It is already doing so. The question is whether legal frameworks will catch up before the transformation is complete—and the evidence suggests they will not, unless urgent action is taken now.

The Convergence That Changes Everything

Three technological trajectories are converging in the 2020s, each accelerating the others. Artificial intelligence, once confined to processing text and images, now decodes neural signals with increasing accuracy. Big data infrastructure, built to harvest and monetize personal information, stands ready to absorb the most intimate data of all—brain activity. And neurotechnology itself, long confined to medical clinics, is entering the consumer market through wireless earbuds, headbands, and soon, implantable chips.

As the United Nations Human Rights Council Advisory Committee documented in its 2024 report (A/HRC/57/61), this convergence creates a situation without precedent: "Neurotechnologies are unique and socially disruptive because they generally: (a) enable the exposition of cognitive processes; (b) enable the direct alteration of a person's mental processes and thoughts; (c) bypass the individual's conscious control or awareness; (d) enable non-consensual external access to thoughts, emotions and mental states."

Each of these capabilities is alarming on its own. Together, they represent a fundamental challenge to the architecture of human rights, which was designed to protect the physical person and their external expressions—not the inner workings of the mind itself.

The Market Is Moving Faster Than the Law

The economic pressure is immense and growing. According to the UN report, private investment in neurotechnology has increased "from $331 million to $7.3 billion in just 10 years." Major technology companies are racing to place products in the market, and "nearly half of all worldwide patent applications are registered in the United States."

Yet this explosive growth is occurring in what the report calls "a highly deregulated environment." Consumer neurotechnology devices—headbands that claim to improve concentration, earbuds that monitor brain activity, gaming interfaces that respond to mental states—face none of the regulatory scrutiny applied to medical devices. "In contrast to the medical development of neurotechnologies, which are generally regulated by established frameworks and processes, the consumer sector remains highly deregulated and is characterized by rapid expansion."

The result is a regulatory vacuum that companies are actively filling. As the UN report warns: "Existing loopholes in regulations, lack of technical expertise and capacity and the absence of adequate oversight bodies are factors that will certainly be exploited by large companies seeking profits."

What Is Already Possible

The technologies that raise these concerns are not speculative. They exist now. Non-invasive neurotechnology devices are "used for a variety of medical purposes, but the same devices are also being commercialized for different applications in the area of well-being, in the workplace or as an educational tool." Users can "wirelessly track their levels of concentration and check their results on their phone."

Meanwhile, invasive technologies are advancing rapidly. One company "has recently developed a secure interface for communication 'with the power of thought'" and is "conducting large-scale trials of this technology, which can be implanted into the brain through blood vessels." Other companies "already advertise 'cosmetically' invisible implants that may allow users to control computers or mobile devices from any location" and "claim to be working on the 'cognitive enhancement, pleasure or leisure of healthy people.'"

The UN report is careful to note that "most neurotechnologies are still being tested and do not provide accurate results in the decoding of people's thoughts, which means that 'mind-reading' or 'mind-steering' are not still achievable." But this is cold comfort. Neuroimaging can already "be employed to make inferences about individuals on a large spectrum of mental states, including memories, semantic knowledge, emotions, dreams, inner speech and intentions." And with AI assistance, these inferences grow more accurate with each iteration.

The Rights at Immediate Risk

The UN report identifies multiple rights that are already threatened. Freedom of thought, long considered an abstract provision, now demands concrete definition. "Since neurotechnologies have the ability to interfere with the human mind, freedom of thought comes to the fore." The report emphasizes that this right "does not allow for any interference whatsoever and must be protected unconditionally."

Privacy faces new and unprecedented threats. "Neurodata provide highly sensitive information and allow aspects of individuals' identity and personality to be revealed, such as sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions." The report warns that "consumer neurotechnology companies are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future."

Personal integrity—the right to control one's own body and mind—is directly implicated. "Mental integrity would also be affected in cases of unauthorized access to mental activity with the aim of intruding into mental processes, altering mental functioning or seeking to manipulate."

And the prohibition on torture and cruel treatment, which the report calls "absolute," applies to neurotechnologies that "constitute an inherently degrading treatment in all circumstances," particularly "brain-reading and brain-writing techniques."

The Window Is Closing

The UN report's central insight is that "when dealing with highly disruptive technologies anticipation is key." It warns that "human rights violations may happen as a consequence of actions and misuses, but they are also the result of lack of preparedness."

We are in a moment of preparedness—or the last chance to achieve it. The technologies are advancing, the market is growing, the legal frameworks are absent, and the window for establishing protections before widespread deployment is closing rapidly. "The risk is that, without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations."

This is why now. Not next year, not when the first major scandal breaks, not when the first implantable consumer device reaches millions of users. Now—while the technologies are still imperfect, while the legal landscape is still blank, while there is still time to build the protections that the inner realm of thought requires.

PART II

The International Legal Landscape

The year 2024 marks a watershed moment for neuro-rights, not because the legal frameworks are complete—they are not—but because the international community has finally begun to articulate the terms of the debate. Across multiple jurisdictions and institutional settings, a patchwork of initiatives is emerging, each addressing a different facet of the challenge posed by neurotechnology. Together, they form the foundation of what may become a global legal architecture for the protection of the human mind.

A Decade of Neuro-rights — Milestones 2017 – 2026

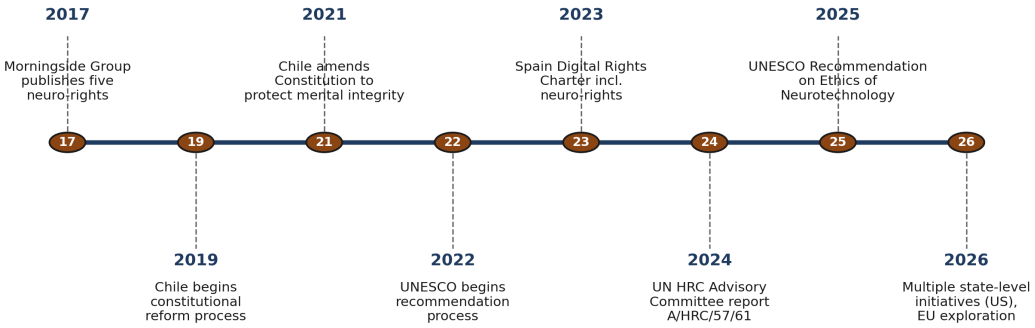


Figure 6.1. *A decade of neuro-rights — from the Morningside Group's 2017 proposal of the five rights to the UNESCO Recommendation of 2025 and the multiple state-level initiatives now under way.*

The United Nations Takes the Lead

The most significant development at the international level is the United Nations Human Rights Council resolution 51/3, adopted on 6 October 2022. Co-sponsored by 23 states, this landmark resolution requested the Advisory Committee to prepare a comprehensive study on "the impact, opportunities and challenges of neurotechnology with regard to the promotion and protection of all human rights." The resulting report, A/HRC/57/61, published on 8 August 2024, is the most authoritative international examination of neurotechnology's human rights implications to date.

The report's drafting group, established in February 2023, brought together experts from diverse regions and legal traditions. Their work represents the first systematic attempt by an international human rights body to map the terrain. The report identifies six rights "particularly at risk": freedom of thought, the right to privacy, the right to personal integrity, the right to health, the prohibition of torture, and the right to a fair trial. It does not propose a new treaty or binding instrument, but it provides the analytical foundation upon which such instruments could be built.

Chile's Constitutional Precedent

Before the UN report, one country had already acted. In 2021, Chile became the first nation to amend its constitution to protect brain rights. The reform added a provision to Article 19, guaranteeing "the protection of the brain and its mental activity" against unauthorized interference. While the amendment is brief and its implementation remains contested, its symbolic importance cannot be overstated. Chile demonstrated that constitutional protection for the inner realm of thought is not a theoretical abstraction but a practical possibility.

Regional Frameworks in Development

The European Union's AI Act, adopted in 2024, represents the most ambitious attempt to regulate artificial intelligence, but its coverage of neurotechnology is incomplete. The Act classifies certain AI systems used in neurotechnology as "high-risk," triggering requirements for risk assessment, transparency, and human oversight. However, as critics have noted, the Act fails to address implants, directed energy devices, and non-invasive mind-reading technologies that operate without AI integration. The result is a regulatory framework with significant gaps.

The Council of Europe's Framework Convention on AI, also adopted in 2024, takes a more principled approach, grounding its provisions in existing human rights obligations. Yet it too lacks specific provisions for neurotechnology, leaving member states to determine how its general principles apply to brain-computer interfaces and neural data.

State-Level Action in the United States

In the absence of federal legislation, two U.S. states have moved independently. Colorado, in 2024, amended its privacy law to classify neural data as "sensitive data," requiring explicit consent for collection and processing. California followed with a similar amendment to its Consumer Privacy Act. These state-level actions create a patchwork of protections but also establish precedents that federal legislation could adopt.

The Challenge of Coordination

The international legal landscape is thus characterized by fragmentation. Chile's constitutional amendment exists alongside the UN's non-binding report, the EU's partial regulatory framework, and U.S. state-level privacy laws. No single instrument provides comprehensive protection. No treaty addresses neuroweapons. No international body has authority to enforce neuro-rights across borders.

The UN report itself acknowledges this gap: "Developing context-specific standards may be a necessary measure to better define the scope of rights being engaged." It calls for "reinforcing the institutional structure to ensure that prevention and mitigation policies are effectively implemented." The question is whether the international community can move from analysis to action before the technologies outpace the law.

CHAPTER 8

Chile's Pioneering Constitutional Reform (2021)

In October 2021, Chile became the first country in the world to amend its constitution to protect brain rights, establishing a legal precedent that would reverberate through international human rights discourse. The reform, driven by Senator Guido Girardi and supported by a coalition of neuroscientists, legal scholars, and human rights advocates, inserted a new provision into Article 19 of the Chilean Constitution, declaring that technological developments must respect the physical and mental integrity of individuals.

The amendment's language was deliberately broad, recognizing that neurotechnology would evolve faster than legislation could track. It established that "scientific and technological development must be at the service of people and carried out with respect for life and physical and mental integrity." More specifically, it mandated that the law regulate the requirements, conditions, and restrictions for the use of neurotechnology, ensuring the protection of brain activity and the information derived from it.

The Political Context

Chile's constitutional reform emerged from a unique political moment. The country was in the midst of rewriting its entire constitution following the 2019 social protests, creating a rare opportunity for fundamental rights innovation. Senator Girardi, a medical doctor by training, had been tracking neurotechnology developments closely, particularly the work of Columbia University neuroscientist Rafael Yuste, who had been advocating for "neurorights" since 2017.

The reform process moved with unusual speed for constitutional change. The proposal passed through both chambers of Chile's Congress with broad bipartisan support, reflecting a rare consensus that the protection of mental privacy required immediate action. The final text was promulgated on October 25, 2021, and published in the official gazette on November 2, 2021.

What the Amendment Actually Does

The Chilean amendment operates at two levels. First, it establishes a constitutional principle: that brain activity and neural data deserve special protection because they constitute the "last private place" of human identity. This principle creates a baseline that all subsequent legislation must respect. Second, it mandates the development of specific laws to implement this protection, covering areas such as:

Limits on commercial use: Prohibiting the collection, storage, or sale of neural data without explicit, informed consent.

Prohibition of manipulation: Banning the use of neurotechnology to alter mental processes without the individual's knowledge or consent.

Protection of personal identity: Ensuring that neurotechnology cannot be used to undermine an individual's sense of self or agency.

Guarantee of free will: Protecting the capacity for autonomous decision-making from technological interference.

Implementation Challenges

The constitutional amendment was a necessary first step, but implementation has proven challenging. Chile's Congress has yet to pass the enabling legislation that would give the amendment practical force. Critics note that the constitutional language is aspirational rather than operational, leaving significant gaps in enforcement.

The UN Human Rights Council report (A/HRC/57/61) implicitly acknowledges this challenge, noting that "developing context-specific standards may be a necessary measure to better define the scope of rights being engaged." Chile's experience demonstrates that constitutional protection, while symbolically powerful, requires accompanying legislative and institutional frameworks to be effective.

International Impact

Despite its implementation challenges, Chile's reform has had significant international influence. It provided a template for subsequent neuro-rights initiatives in other countries, including Brazil, Mexico, and Spain, which have introduced similar constitutional or legislative proposals. The reform also informed the UN Human Rights Council resolution 51/3 (2022), which requested the study that would become the 2024 UN report on neurotechnology and human rights.

The Chilean precedent established that brain rights could be framed as an extension of existing human rights—particularly the right to mental integrity and privacy—rather than requiring an entirely new rights framework. This approach has been adopted by subsequent initiatives, including the Colorado and California neurodata protection laws of 2024.

The Unfinished Agenda

Chile's constitutional reform remains a pioneering but incomplete achievement. The amendment created a constitutional floor for neuro-rights, but the ceiling—the specific legislation needed to operationalize those rights—has yet to be built. The UN report's warning that "existing loopholes in regulations, lack of technical expertise and capacity and the absence of adequate oversight bodies are factors that will certainly be exploited" applies directly to Chile's unfinished implementation.

The reform's legacy may ultimately lie less in its specific provisions than in its demonstration that constitutional protection of brain rights is politically achievable. As the UN report notes, "anticipation is key" when dealing with disruptive technologies. Chile acted before the technologies matured, establishing a principle that other nations could follow. Whether that principle translates into effective protection will depend on the legislative and institutional steps that follow.

CHAPTER 9

The UN Human Rights Council Resolution 51/3 (2022)

On 6 October 2022, the United Nations Human Rights Council took a step that would fundamentally alter the trajectory of neuro-rights governance worldwide. By adopting resolution 51/3, the Council formally recognized that neurotechnology posed questions of such profound human rights significance that they demanded systematic international attention. The resolution requested the Human Rights Council Advisory Committee to prepare a comprehensive study on "the impact, opportunities and challenges of neurotechnology with regard to the promotion and protection of all human rights."

The resolution was co-sponsored by 23 states, reflecting an unusually broad consensus on the urgency of the issue. This coalition included countries from every region, signaling that concerns about neurotechnology were not confined to any single political or geographic bloc. The breadth of support demonstrated that the international community recognized neurotechnology as a genuinely global challenge—one that could neither be addressed by individual nations alone nor deferred to some future date when the technology had fully matured.

The Mandate

The resolution's request was deliberately broad. It asked the Advisory Committee not merely to catalog existing neurotechnologies, but to examine their impact across the full spectrum of human rights—civil, political, economic, social, and cultural. The study was to include "action-oriented recommendations" on how the Human Rights Council, its special procedures, and its subsidiary bodies could address neurotechnology's implications "in a coherent, holistic and inclusive manner."

This mandate reflected a recognition that neurotechnology did not fit neatly into existing human rights categories. It was not solely a privacy issue, nor a health issue, nor a freedom-of-thought issue—though it implicated all of these. The resolution implicitly acknowledged that the technology's disruptive potential required a comprehensive, cross-cutting response.

The Drafting Process

In February 2023, the Advisory Committee established a drafting group to produce the requested study. The group was composed of 14 experts from diverse legal and geographic backgrounds: Nurah Alamro, Noor al-Jehani, Joseph Gérard Angoh, Buhm-Suk Baek, Milena Costas Trascasas (Rapporteur), Riva Ganguly Das, Jewel Major, Javier Palummo, Vasilka Sancin, Patrycja Sasnal, Vassilis Tzevelekos (Chair), Catherine Van de Heyning, Frans Viljoen, and Yue Zhang.

The drafting group's composition was notable for its inclusion of experts from countries that had already engaged with neuro-rights domestically, as well as representatives from

regions—such as the Global South—that might be disproportionately affected by neurotechnology's development. This diversity was intentional. The study was to be not merely a technical assessment but a document that reflected the full range of human rights concerns raised by neurotechnology across different legal traditions and levels of economic development.

The Context

Resolution 51/3 did not emerge from a vacuum. It was preceded by years of growing awareness among human rights bodies that neurotechnology represented a unique category of disruptive innovation. The UN report that would eventually result from this resolution noted that "for a long time, human mind and conscience were considered the container of the inner self, a fortress, inaccessible from the outside." Neurotechnology threatened to breach that fortress.

The resolution also reflected a broader shift in the international community's approach to emerging technologies. Rather than waiting for violations to occur and then responding, the Human Rights Council chose to act proactively. As the subsequent UN report would emphasize, "when dealing with highly disruptive technologies anticipation is key." Resolution 51/3 embodied this principle: it sought to understand the risks before they materialized, to build guardrails before the technology became ubiquitous.

The Significance

Resolution 51/3 was significant for several reasons. First, it placed neurotechnology squarely within the human rights framework, rejecting the notion that these were merely ethical or technical questions. Second, it acknowledged that existing human rights protections might be insufficient to address the specific challenges posed by direct access to the human mind. Third, it set in motion a process that would culminate in the comprehensive 2024 UN report (A/HRC/57/61), which would provide the most detailed international analysis yet of neurotechnology's human rights implications.

The resolution also established a precedent for future international action. By requesting a study with "action-oriented recommendations," the Human Rights Council signaled that it was not merely interested in academic analysis but in concrete steps that could be taken to protect human rights in the neurotechnology age. This orientation would influence subsequent initiatives, including UNESCO's 2025 Recommendation on the Ethics of Neurotechnology and the ongoing discussions about a potential international treaty on neuro-rights.

In the two years between the resolution and the publication of the requested study, the neurotechnology industry would grow from $331 million in investments to $7.3 billion—a more than twenty-fold increase. The urgency that the resolution had identified would only intensify. But the foundation had been laid. The international community had formally recognized that the battle for the last private place required a coordinated, global response.

CHAPTER 10

The 2024 UN Report on Neurotechnology and Human Rights

On 8 August 2024, the United Nations Human Rights Council published a document that would become the most comprehensive international analysis of neurotechnology's human rights implications ever produced. The report, formally designated A/HRC/57/61 and titled *Impact, Opportunities and Challenges of Neurotechnology with regard to the Promotion and Protection of All Human Rights*, was the direct result of the process set in motion by Resolution 51/3 two years earlier.

The Mandate and Process

The report was prepared by the Human Rights Council Advisory Committee, a body of independent experts tasked with providing the Council with studies and research-based advice. In February 2023, the Committee established a drafting group composed of fourteen members from diverse regions: Nurah Alamro, Noor al-Jehani, Joseph Gérard Angoh, Buhm-Suk Baek, Milena Costas Trascasas (who served as Rapporteur), Riva Ganguly Das, Jewel Major, Javier Palummo, Vasilka Sancin, Patrycja Sasnal, Vassilis Tzevelekos (Chair), Catherine Van de Heyning, Frans Viljoen, and Yue Zhang. The group was charged not merely with describing the technology, but with producing "action-oriented recommendations on how the impact of neurotechnologies on human rights can be addressed within the Human Rights Council, its special procedures and its subsidiary bodies in a coherent, holistic and inclusive manner."

A Unique Disruptive Technology

The report began by establishing why neurotechnology demanded special attention. While acknowledging the "high expectations" raised by rapid progress in the field—including new tools for diagnosis, treatment of neurological and psychiatric diseases, and assistive devices that could improve mobility and autonomy—the report identified six characteristics that made neurotechnologies "unique and socially disruptive":

First, they enable the exposition of cognitive processes. Second, they enable the direct alteration of a person's mental processes and thoughts. Third, they can bypass the individual's conscious control or awareness. Fourth, they enable non-consensual external access to thoughts, emotions, and mental states. Fifth, they are nurtured by "neurodata"—information gathered from the brain and nervous system, including inferences extracted from it—which are needed for their own functioning, calibration, and optimization. Sixth, they collect, analyze, and process large personal datasets of a highly sensitive nature.

The report was careful to distinguish between current capabilities and science fiction. "Despite their rapid progress," it noted, "most neurotechnologies are still being tested and do not provide accurate results in the decoding of people's thoughts, which means that 'mind-reading' or 'mind-steering' are not still achievable." However, neuroimaging was

"increasingly explored and can be employed to make inferences about individuals on a large spectrum of mental states, including memories, semantic knowledge, emotions, dreams, inner speech and intentions." The convergence with artificial intelligence and machine learning, the report warned, would exacerbate issues such as algorithmic bias, lack of traceability, and lack of transparency.

The Deregulated Industry

The report painted a stark picture of the commercial landscape. Large-scale "brain initiatives" had been launched by multiple countries since 2013—the United States, Australia, China, the European Union, Israel, Japan, and the Republic of Korea. Public funding and private investments had fueled an industry that grew from $331 million to $7.3 billion in just ten years. Nearly half of all worldwide patent applications were registered in the United States.

The growing role of the private sector was changing the dynamics of the industry, leaning it toward direct-to-consumer devices. "With breakthrough advances reported daily," the report noted, "the pressure from large technology companies to place neurotechnology products in the market will inevitably increase." Yet "in contrast to the medical development of neurotechnologies, which are generally regulated by established frameworks and processes, the consumer sector remains highly deregulated and is characterized by rapid expansion."

The report warned that "existing loopholes in regulations, lack of technical expertise and capacity and the absence of adequate oversight bodies are factors that will certainly be exploited by large companies seeking profits." The risk was that "without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations."

Non-invasive devices were already on the market—headsets, wristbands, electrode patches, tattoos, earphones—used to improve cognitive and sporting performance or to aid concentration, relaxation, or sleep. The new generation of wireless earbuds included biosensing electrodes enabling measurement of brain activity. Yet the report noted that "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge" and that "research on potential side-effects, long-term impacts and their possible irreversible character is lacking."

The Human Rights Framework

The core of the report examined how neurotechnologies engaged specific human rights. It began with the foundational principle of human dignity, which "encapsulates the essence of what it means to be human." The report insisted that dignity required assessing neurotechnologies' capacity to affect "the constitutive elements of the person: the individual's personhood, mental capacities and personality."

Freedom of thought received the most extensive treatment. The report reaffirmed that freedom of thought is an absolute right—one that "does not allow for any interference whatsoever." It protected the *forum internum*, "the inner, psychological realm of the person,

where thoughts and convictions are formed." Neurotechnologies could serve as "mindreading-like devices" to make inferences about individuals, and with AI support could "elaborate sophisticated individualized psychological profiles and use them to predict individuals' behaviour or intentions or to infer their mental characteristics or predispositions." The report warned that neuromodulation devices could alter mental processes "in ways similar to 'brainwashing'" and that neuromarketing strategies could "condition the forming of opinions, as well as influencing an individual's decision-making processes."

The right to privacy was threatened by neurotechnologies' capacity to collect and record personal information from the mind itself. Neurodata could reveal "sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions." The report noted that "consumer neurotechnology companies are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future."

The right to personal integrity was engaged when neurointerventions "produce effects on the body and the psychological sphere." The report warned that some technologies could "affect mental health and provoke alterations in an individual's personality, psychological balance or sense of self-identity." It emphasized that "enhanced mechanisms should be adopted to ensure that persons in vulnerable situations are effectively protected from all forms of coercion."

The right to health was both a promise and a risk. While neurotechnology held "great prospects in protecting mental health," the report cautioned that "in the absence of a medical need and without being supported by extensive trials and scientific evidence, devices may become harmful for mental health and their extensive use may lead to public health problems."

The prohibition of torture and cruel, inhuman or degrading treatment was absolute. The report stated that "brain-reading and brain-writing techniques, particularly 'neurocorrectives,' constitute an inherently degrading treatment in all circumstances." It affirmed that "no statements or confessions or other evidence obtained in violation of article 7 of the International Covenant on Civil and Political Rights may be invoked in any proceedings, not even during a state of emergency."

The right to a fair trial was threatened by proposals to use neurotechnologies in criminal investigation. The report found it "difficult to see how they can be used without violating the right to due process of law," noting that "subjecting detained or imprisoned persons to 'brain-reading' techniques is clearly a form of prohibited pressure."

Vulnerable Groups

The report identified persons with disabilities as particularly exposed to risk, noting that neurotechnologies were "often developed without proper human rights guidelines and monitoring." It insisted that "countering stereotypes and ableism requires prioritizing the needs, preferences and rights of persons with disabilities" and that their representative organizations "should be involved in the development of products that directly affect them."

The Anticipatory Approach

The report's central methodological contribution was its insistence on anticipation. "When dealing with highly disruptive technologies anticipation is key," it stated. "Although human rights violations may happen as a consequence of actions and misuses, they are also the result of lack of preparedness." The report called for "developing specific policies, legislation and instruments" and "reinforcing the institutional structure to ensure that prevention and mitigation policies are effectively implemented."

The 2024 UN report did not merely describe the problem—it provided a framework for action. It clarified that existing human rights instruments, while foundational, required context-specific elaboration to address the unique challenges of neurotechnology. It identified the rights most at risk and the groups most vulnerable. And it established that the international community could not afford to wait for violations to occur before responding. The report was both a warning and a roadmap—a document that would shape the neuro-rights debates for years to come.

CHAPTER 11

The 2025 UNESCO Recommendation on the Ethics of Neurotechnology

If the 2024 UN report was the diagnosis, the 2025 UNESCO Recommendation on the Ethics of Neurotechnology was the prescription. Adopted by UNESCO's General Conference in November 2025, this landmark instrument represented the first global ethical framework specifically designed to govern neurotechnology. It built directly upon the foundation laid by the UN Human Rights Council's resolution 51/3 and the 2024 report, translating human rights principles into actionable ethical standards for governments, researchers, and industry.

The Context: Why UNESCO?

The choice of UNESCO as the forum for this framework was deliberate. Since its founding, UNESCO had served as the global platform for ethical debates on emerging technologies—from the Universal Declaration on the Human Genome and Human Rights (1997) to the Recommendation on the Ethics of Artificial Intelligence (2021). Neurotechnology, which directly engages questions of human identity, consciousness, and dignity, naturally fell within UNESCO's mandate.

The process began in earnest following the 2024 UN report. Member states recognized that while the Human Rights Council could identify rights at risk, a more detailed ethical framework was needed to guide national legislation and international cooperation. UNESCO's Intergovernmental Bioethics Committee and World Commission on the Ethics of Scientific Knowledge and Technology (COMEST) had already produced preliminary work, including the 2022 *Ethical Issues of Neurotechnology* report. The 2025 Recommendation consolidated these efforts into a binding normative instrument—not a treaty, but a recommendation that member states were expected to implement through national legislation.

Core Principles of the Recommendation

The Recommendation was organized around four core ethical principles, each elaborated with specific provisions:

1. Human Dignity and the Primacy of the Person. The Recommendation placed human dignity at the center, echoing the 2024 UN report's insistence that neurotechnology must never reduce individuals to objects of manipulation. It affirmed that "the human person shall be the primary subject of neurotechnology development and application, not a mere means to scientific or commercial ends." This principle explicitly prohibited applications that "undermine the constitutive elements of personhood, including mental capacities, personality, and self-determination."

2. Autonomy and Informed Consent. The Recommendation recognized that neurotechnology posed unique challenges to consent. Unlike traditional medical procedures, neurotechnology could bypass conscious reasoning, leaving individuals "without psychological protection from having their mind involuntarily read." The framework required that consent be "free, prior, informed, and revocable at any time without detriment to the individual." It further stipulated that consent must be based on "comprehensive understanding of the technology's capabilities, limitations, and potential long-term effects"—a standard that effectively prohibited many direct-to-consumer devices whose risks remained unknown.

3. Privacy and Data Protection. Building on the 2024 report's identification of neurodata as "highly sensitive information," the Recommendation called for "specific legal protections for neural data that recognize its unique character as biometric information capable of revealing aspects of identity, personality, and mental states." It required that neurodata be collected only for "specified, explicit, and legitimate purposes" and prohibited its "repurposing, sale, or transfer to third parties without explicit and renewed consent." The framework also addressed the challenge of algorithmic inference, recognizing that AI-powered analysis could extract information far beyond what raw neural signals revealed.

4. Justice and Equity. The Recommendation confronted the emerging "neuronal divide" head-on. It called on states to "ensure equitable access to neurotechnology for therapeutic purposes, particularly for persons with disabilities and those in underserved communities." At the same time, it warned against "unregulated cognitive enhancement that could exacerbate existing inequalities and create new forms of discrimination." The framework explicitly prohibited "the use of neurotechnology to discriminate against individuals on the basis of neural characteristics or cognitive profiles."

Specific Prohibitions and Safeguards

The Recommendation went beyond general principles to establish specific prohibitions:

Non-consensual neuromodulation was prohibited in all circumstances, with the Recommendation stating that "no person shall be subjected to neurointerventions that alter their mental processes without their free and informed consent."

Coercive uses in military and law enforcement were restricted, with the framework calling on states to "refrain from developing or deploying neurotechnology for purposes of interrogation, behavior modification, or cognitive warfare."

Surveillance applications were limited, particularly in workplaces and educational settings, where the Recommendation required "strict necessity and proportionality assessments" before any neurotechnology-based monitoring could be implemented.

Commercial exploitation of neurodata was constrained, with the framework insisting that "neural data shall not be treated as a commodity subject to unrestricted commercial exploitation."

Implementation Mechanisms

The Recommendation did not merely state principles—it provided a roadmap for implementation. It called on member states to:

7. **Establish national ethics committees** with expertise in neuroscience, human rights, and ethics to oversee neurotechnology development and deployment.
8. **Conduct human rights impact assessments** before approving new neurotechnology applications, particularly those intended for non-medical use.
9. **Develop regulatory frameworks** that address both medical and consumer neurotechnology, closing the regulatory gaps that had allowed the direct-to-consumer market to expand with inadequate oversight.
10. **Promote public awareness and education** about neurotechnology's capabilities, risks, and rights implications.
11. **Foster international cooperation** through information sharing, joint research on safety and ethics, and mutual recognition of regulatory standards.

The Significance of the Recommendation

The 2025 UNESCO Recommendation was not legally binding in the same sense as a treaty, but its normative weight was substantial. As a product of UNESCO's standard-setting process, it carried the authority of the international community's collective judgment on ethical neurotechnology governance. Member states were expected to report on their implementation progress, and the framework provided a benchmark against which national legislation could be evaluated.

Critically, the Recommendation filled a gap that the 2024 UN report had identified: the need for "context-specific standards" that could "better define the scope of rights being engaged." By translating human rights principles into specific ethical guidelines, UNESCO provided legislators, regulators, and industry with a practical tool for navigating the neurotechnology landscape.

The Recommendation also served as a foundation for future binding instruments. Just as the Universal Declaration of Human Rights had preceded binding covenants, the UNESCO Recommendation established the normative groundwork for what many advocates hoped would become an international treaty on neuro-rights. It was, in the words of one delegate, "not the end of the conversation, but the beginning of a new phase—one in which ethical principles become legal obligations."

CHAPTER 12

Spain, Brazil, and the Latin American Neuro-rights Movement

Chile's 2021 constitutional reform did not occur in isolation. It was the first crystallization of a regional intellectual movement that has spread, with varying intensity, across Spanish-speaking and Portuguese-speaking jurisdictions. Latin America is, by some distance, the part of the world where neuro-rights have moved fastest from academic proposal to legislative draft. Spain joined the conversation as the Iberian bridge between Latin American activism and European Union policymaking. Brazil and Mexico have produced their own draft bills. Smaller initiatives are underway in Argentina, Colombia, Uruguay, and Peru. Understanding this regional cluster is essential to predicting where the next constitutional or statutory breakthroughs are most likely to occur.

Why Latin America Moved First

Several factors combine to explain the regional lead. The first is intellectual lineage. Rafael Yuste, the Spanish-born Columbia University neuroscientist who has done more than any other individual to articulate the neuro-rights framework, has direct ties through language and family to the Spanish-speaking world. He has lectured extensively in Chile, Spain, Mexico, and Argentina, and his arguments have circulated in those jurisdictions in their original language rather than in translation. The second factor is the constitutional tradition of much of Latin America, which tends to recognize a broader catalogue of rights than common-law systems and which has shown willingness to incorporate new rights as constitutional questions emerge. The third is the politics of human rights itself. Many Latin American democracies were rebuilt after periods of authoritarian rule in which the state intruded on the inner lives of citizens — through torture, surveillance, and ideological indoctrination. The cultural memory of those abuses makes the protection of the inner forum a more visceral political concern than in jurisdictions whose democratic traditions have not been so recently interrupted.

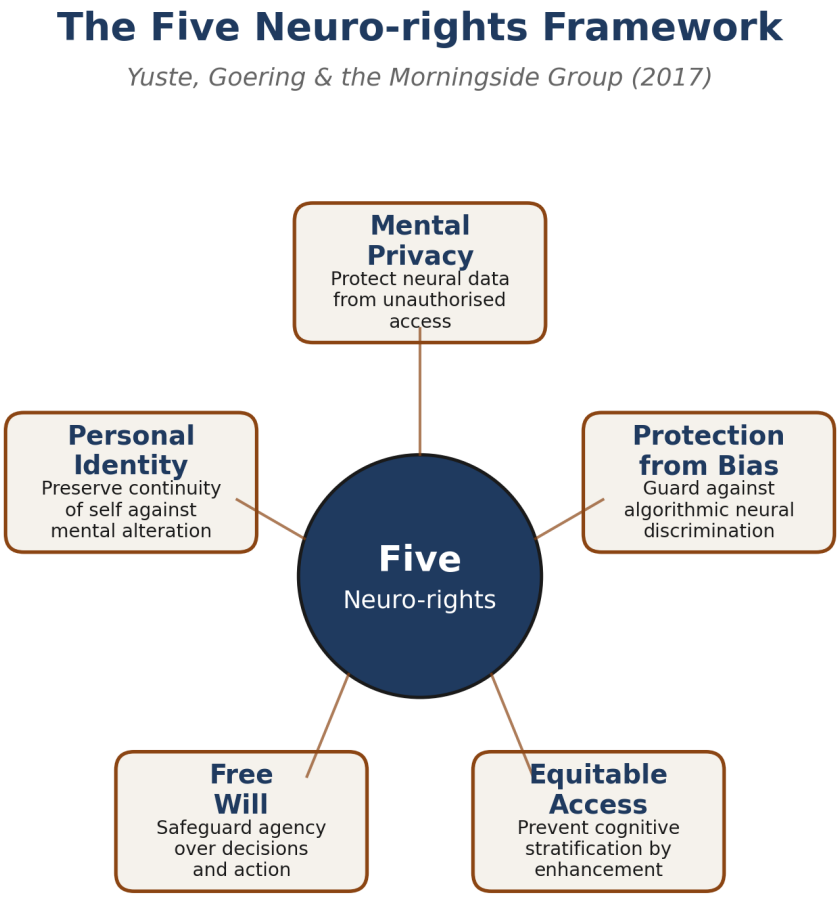


Figure 5.1. *The five neuro-rights framework articulated by Yuste, Goering and the Morningside Group in 2017 — the conceptual scaffold of every neuro-rights instrument adopted since.*

Spain: The Charter of Digital Rights

Spain's contribution to the neuro-rights conversation came not through constitutional amendment but through the 2021 Charter of Digital Rights (Carta de Derechos Digitales), promulgated by the government of Pedro Sánchez. The Charter is, in formal legal terms, a soft-law instrument: it does not create directly enforceable individual rights but commits the executive to interpret existing law and craft future legislation consistent with its principles. Section XXVI of the Charter — Digital Rights in the Use of Neurotechnologies — sets out four principles in language that has since echoed through every subsequent European draft.

First, the Charter affirms that the conditions, limits and guarantees of the implementation and use of neurotechnologies must ensure protection of human dignity, individual autonomy and self-determination, mental privacy, the right not to be subjected to interventions that limit cognitive identity, and the right to control decisions taken by automated processes that interact with brain activity. Second, it specifies that any treatment of brain data requires explicit, free, and informed consent and must be subject to particular protective measures.

Third, it asserts a right against intervention that would alter cognitive identity. Fourth, it calls for transparency about the algorithms involved in any system that processes brain data.

What makes the Spanish Charter important is less its content — broadly similar to the Chilean amendment — than its political position. As an executive instrument of an EU Member State, the Charter signals to Brussels that at least one large Member State considers neurotechnology to require dedicated protection. It has been cited in European Parliament debates, in legal commentary on the EU AI Act, and in successive draft proposals for European-level neuro-rights frameworks. Without the Charter, the European conversation would lag the Latin American one by years.

Brazil: The Senate Bill

Brazil's contribution has taken the form of a federal legislative bill, PL 4.690 of 2023, introduced in the Senate by Senator Randolfe Rodrigues. The bill proposes amendments to the Brazilian Constitution to recognize neuro-rights and to the General Data Protection Law (LGPD) to classify neural data as a special category of sensitive personal data requiring heightened protection. As of writing, the bill is at committee stage. Its passage is uncertain. Its mere existence, however, has shifted the legal centre of gravity in the largest country in South America.

The Brazilian bill is structurally interesting because it pairs constitutional amendment with data-protection statute. The constitutional language secures the framework. The statutory language operationalizes it. Chile's experience demonstrated the limits of constitutional reform without accompanying legislation. Brazil's drafters appear to have learned that lesson. If the bill passes in something close to its current form, Brazil will become the second country with explicit neuro-rights in its constitution and the first to combine that constitutional protection with a fully developed statutory regime.

Mexico and the Federal Initiative

Mexico has seen multiple legislative initiatives related to neuro-rights, the most substantive of which is a proposed amendment to Article 4 of the Constitution that would recognize a right to mental integrity and a right to access information generated from one's own neural data. The Mexican proposal is notable for emphasizing the data-access right: a citizen has a right to know what neural data has been collected about them and to receive a copy of that data on request. This emphasis mirrors elements of European data-protection law and represents a useful synthesis of the regional debate with international privacy norms.

Mexico's proposal has not yet passed. Its progress is conditioned by the broader political dynamics of the Mexican congress, which has had other priorities. But the proposal is being debated, has support from neuroscientists and human-rights organizations, and represents a serious legislative effort in a country whose constitutional decisions reverberate through Central and parts of South America.

Argentina, Colombia, and the Smaller Initiatives

Beyond the three leading jurisdictions, several other Latin American countries have neuro-rights initiatives at various stages of development. Argentina has seen proposals in the Buenos Aires provincial legislature. Colombia has a pending initiative in the federal congress. Uruguay's executive has commissioned a study on neurotechnology and human rights. Peru has produced academic commentary that points toward future legislative drafts. Each of these initiatives is at an earlier stage than Brazil's or Mexico's, but together they create a critical mass: a regional cluster of activity that signals to international bodies that the Latin American position on neuro-rights is settled and that any global treaty negotiation will have a substantial pro-protection bloc from the start.

The Inter-American Declaration of Principles

An important regional document is the Inter-American Declaration of Principles Regarding Neuroscience, Neurotechnologies, and Human Rights, produced under the auspices of the Inter-American Juridical Committee of the Organization of American States. The Declaration is a soft-law instrument that articulates principles for the use of neuroscience and neurotechnology consistent with the Inter-American human-rights framework. It does not bind member states, but it provides interpretive guidance for the Inter-American Commission and Court of Human Rights and serves as a template for national legislation across the hemisphere.

The Declaration's principles include affirmations of mental privacy, free will, personal identity, equitable access to cognitive enhancement, and protection from bias. It also articulates principles concerning the use of neurotechnology in criminal justice and in armed conflict. The Declaration's framing of these issues is more explicit than the corresponding language in the UN Human Rights Council report and is, in places, more protective. Citizens and legislators in OAS member states should be aware that an interpretive instrument exists and that domestic courts can be invited to read national law consistently with it.

Country	Instrument	Status	Key Innovation
Chile	Constitutional amendment to Art. 19	Adopted (2021)	First constitutional recognition; mental integrity as right
Spain	Carta de Derechos Digitales § XXVI	Adopted (2021), soft law	Executive interpretive instrument; influence in EU debate
Brazil	PL 4.690/2023 + LGPD amendment	Committee stage	Pairs constitutional and statutory reform
Mexico	Amendment to Article 4	Pending	Strong access right to one's neural data
Argentina	Provincial proposals (Buenos Aires)	Drafting	Provincial-level experimentation

Country	Instrument	Status	Key Innovation
Colombia	Federal initiative	Pending	Linkage to mental-health rights
Uruguay	Executive study	Underway	Pre-legislative groundwork
Peru	Academic and civil-society proposals	Pre-legislative	Framework for future drafting
OAS	Inter-American Declaration of Principles	Adopted, soft law	Interpretive instrument for hemisphere

What the Latin American Movement Teaches

Three lessons emerge from the Latin American experience that legislators elsewhere should internalize. The first is that constitutional protection is necessary but not sufficient. Chile demonstrated that an amendment without enabling legislation is symbolically powerful but practically limited. Brazil's drafters have responded by pairing constitutional and statutory work. Future jurisdictions should plan from the start to enact both.

The second lesson is that the regional intellectual leadership of one or two figures can dramatically accelerate adoption. Rafael Yuste's work in Spanish made the conversation accessible to Latin American legislators years before the equivalent English-language conversation reached comparable depth. Other regions — South Asia, the Arab world, sub-Saharan Africa — lack equivalent figures and equivalent linguistic accessibility. Building that capacity through translation, fellowships, and regional research centres is a precondition for taking the conversation global.

The third lesson is that the political case for neuro-rights resonates most strongly where the cultural memory of state intrusion on the inner forum is most recent. Countries whose democratic transitions are still within living memory — Chile, Brazil, Spain, Argentina — have been receptive to neuro-rights arguments in ways that more established democracies have not. As neurotechnology spreads, the rest of the world will encounter the question whether to wait until abuse generates the necessary political memory, or to act on the inheritance of those who do remember.

CHAPTER 13

China and the Authoritarian Use of Neurotechnology

Neuro-rights debates conducted in liberal democracies tend to assume a worst-case scenario in which technology might, hypothetically, be used to surveil and control citizens. In the People's Republic of China, this scenario is no longer hypothetical. Public statements by Chinese officials and researchers, declassified procurement documents, peer-reviewed publications, and reports from journalists and human-rights organizations describe an integrated programme of neurotechnology research, deployment, and military development that operates without the constraints that frame the Western conversation. China is not the only state with such ambitions, but it is the most advanced, the most public about its goals, and the most consequential for the global trajectory.

Military Brain Science as State Doctrine

The People's Liberation Army has formally identified "Military Brain Science" as a strategic research priority. Publications by PLA-affiliated researchers describe a comprehensive taxonomy of brain functions and corresponding military applications. The taxonomy includes injuring the brain (creating weapons that target cognitive function), interfering with the brain (psychological warfare and cognitive disruption), arming the brain (cognitive enhancement of soldiers), and reading the brain (intelligence gathering and lie detection). Each of these categories is the subject of named research programmes with substantial budgets.

The breadth of this doctrine is significant. Where Western military programmes typically emphasize enhancement of one's own soldiers and defence against adversaries, the Chinese programme explicitly includes offensive applications. The publication of this framing in unclassified Chinese military literature is itself a strategic signal. It tells competing powers that neurotechnology is, in Chinese strategic thinking, on the same footing as the development of nuclear and conventional weapons — a matter for state-directed effort with explicit recognition that the technology will be used.

Brain-Computer Interface Programmes

China hosts several major brain-computer interface research initiatives, both civilian and military. The China Brain Project, launched in the late 2010s and integrated into successive five-year plans, is one of the world's largest neuroscience programmes by funding. It directs resources to fundamental research on brain function, neurological disease, and neural engineering. A parallel set of initiatives within defence-affiliated universities, such as the National University of Defense Technology, focuses specifically on military applications.

Chinese researchers have published peer-reviewed work on BCI applications for fighter-pilot control, drone-swarm management, and integrated soldier-machine teaming. The quality of this research is variable but improving, and the volume is now comparable to the United States and Europe combined. China has also developed close ties between military and

civilian research, blurring the distinction in ways that complicate efforts to apply export controls to dual-use technologies.

The Workplace and the Classroom

Independent journalists have documented the deployment of consumer-grade EEG headsets in Chinese workplaces and classrooms. In factories and corporate offices, devices are sold to employers as productivity-monitoring tools that purport to detect attention, fatigue, and stress levels. In schools, similar devices are marketed as aids for monitoring student focus and engagement. The training data generated by these deployments — millions of hours of EEG recordings labelled by context and behaviour — represents one of the world's largest neural datasets and is, by Chinese law, accessible to state authorities on request.

These deployments have been criticized by Chinese parents, teachers, and journalists, and some specific products have been withdrawn from sale after public criticism. But the underlying practice has not been outlawed. The combination of permissive workplace law, a strong state interest in social surveillance, and a thriving consumer-electronics industry means that the data continues to be generated and to flow toward state-linked actors. There is no equivalent in the West, in any liberal democracy, of the volume of workplace and classroom neural data that China is currently amassing.

Xinjiang and the Question of Mass Surveillance

The most serious allegations concerning Chinese use of neurotechnology relate to the Xinjiang Uyghur Autonomous Region, where credible reports document a comprehensive surveillance infrastructure targeting the Uyghur population. The technologies used are predominantly conventional — cameras, facial recognition, mobile-device tracking, biometric collection — but allegations have emerged regarding the use of neurological screening, lie-detection technologies, and behavioural-assessment tools at scale. The precise nature of these deployments is not fully documented in the public record because Xinjiang is closed to independent journalism and to international human-rights monitors.

What is known is sufficient to raise grave concerns. Detained Uyghurs have reported screening procedures that included EEG-style measurements and other instrumentation. Researchers affiliated with Chinese security ministries have published academic work on the neural correlates of "extremism" and "terrorist intent", frameworks that, in practice, have been applied to broad categories of Uyghur dissent. The use of these frameworks in mass-screening contexts represents one of the most consequential test cases for whether neurotechnology can or will be used at population scale by an authoritarian state.

Strategic Implications

The Chinese approach to neurotechnology has several strategic implications for the global neuro-rights conversation. First, it complicates negotiations for any international treaty. A binding instrument requires consent from major states. China's official position on neuro-rights has been guarded, and its incentives to accept restrictions on military neurotechnology research are limited. The same will be true of other powers, particularly Russia, but China is

the largest and most consequential resister. Negotiations that proceed on the assumption that China can be brought into a comprehensive framework will likely fail; negotiations that proceed on a coalition-of-the-willing basis can succeed but will leave significant gaps.

Second, the Chinese data advantage is structural. The volume of neural data being generated in Chinese workplaces, classrooms, and security contexts exceeds, by orders of magnitude, the volume being generated under comparable consent in Western jurisdictions. To the extent that decoder quality scales with training-set size, Chinese decoders may achieve subject-independent performance faster than Western ones. The strategic implication is that Western jurisdictions that wait too long to engage with this technology may find themselves contesting decisions made by systems trained on data they refused to collect.

Third, the example matters. As other authoritarian-leaning regimes observe Chinese deployments, the technologies, business models, and legal frameworks Chinese actors have developed become available for export. Chinese consumer-EEG products are already exported to dozens of countries. Chinese surveillance technology has been adopted in central Asia, parts of Africa, and elements of the Gulf. The combination of capability and willingness to export creates a vector for the spread of authoritarian-style neurotechnology that is faster than the spread of Western consumer products.

What This Means for Neuro-rights

The Chinese reality should inform, but not dictate, the neuro-rights debate in liberal democracies. It should inform the debate by reminding participants that the technologies under discussion are not hypothetical and that the worst-case applications are already being explored. It should not dictate the debate, in the sense that liberal democracies should not abandon protective frameworks in a misplaced effort to compete with practices they should not emulate. The lesson of the twentieth century is that liberal-democratic societies competed with totalitarian alternatives more successfully when they adhered to their own values than when they imitated their adversaries. The same lesson applies, with renewed urgency, in the neurotechnology age.

Practically, the Chinese situation reinforces three priorities for democratic legislators: protect citizens against domestic abuse, build technical capacity to evaluate foreign technology, and engage in international standard-setting so that Chinese norms do not become global ones by default. The neuro-rights movement is, among other things, a soft-power competition for the future. Whether that competition is won will depend less on the technical capability of any particular national programme than on whether the values it serves are recognizable, defensible, and worth choosing.

CHAPTER 14

The African Union, the Arab World, and Global South Perspectives

The neuro-rights conversation, as conducted in the United Nations and in academic journals, is overwhelmingly dominated by voices from Europe, the Americas, and East Asia. The countries of sub-Saharan Africa, the Arab world, South Asia, and the Pacific have been comparatively silent. This silence is not a signal that the technology is irrelevant in those jurisdictions. It is a signal that the political conditions for engagement have not yet matured. As consumer neurotechnology and the data infrastructures that support it spread — often through Chinese exports and Western multinationals operating with limited local oversight — the Global South will increasingly be a frontier for neurotechnology deployment without local rights frameworks to govern it. Addressing this asymmetry is a central challenge for the global neuro-rights movement.

The African Union and the Malabo Convention

The African Union has produced one major continental instrument relevant to neurotechnology: the African Union Convention on Cyber Security and Personal Data Protection (the Malabo Convention), adopted in 2014. The Convention sets out general principles for the protection of personal data and the security of cyber infrastructure across member states. It does not specifically address neural data or neurotechnology, but its general framework would, in principle, encompass such data as a form of sensitive personal information.

The Convention's limitation is that ratification has been slow. As of writing, only a minority of AU member states have ratified the Malabo Convention, and implementation, where it has occurred, has been uneven. The result is a continental framework on paper without comprehensive operational reality. For neurotechnology specifically, this means that even where governments wish to act, the legal infrastructure on which they would build is often missing the foundational data-protection layer. Building neuro-rights frameworks in Africa, then, is not purely a matter of adding new protections to existing law. It often requires the simultaneous construction of the underlying data-protection law as well.

South Africa and the POPIA

Within the African continent, South Africa stands out for having developed a comprehensive data-protection regime, the Protection of Personal Information Act (POPIA), which became fully operational in 2021. POPIA classifies certain categories of personal information as "special personal information" requiring heightened protection — categories including race, religious belief, biometric information, and health information. Neural data is not explicitly listed but is plausibly covered through the biometric and health categories.

South African scholars and civil-society organizations have begun to engage with neuro-rights, and there are early signs that the Information Regulator may issue guidance on the

application of POPIA to neural data. South Africa's combination of relatively developed law, active civil society, and increasing exposure to consumer neurotechnology makes it the most likely candidate for the first sub-Saharan African neuro-rights initiative. The country's experience will shape what is possible elsewhere on the continent.

The Arab Region: Variations on Sovereignty

The Arab region is heterogeneous, and generalizations are unsafe. But several themes recur across multiple jurisdictions. First, there is intense state interest in technologies that can monitor and influence populations, particularly in the wake of the Arab Spring uprisings of the early 2010s. Second, there is investment, particularly in the Gulf states, in becoming hubs for advanced technology in fields including artificial intelligence and biotechnology. Third, there is engagement with international human-rights instruments that is selective: instruments concerning economic and social rights are often more accepted than those touching civil and political rights.

The United Arab Emirates and Saudi Arabia have both made substantial public investments in artificial-intelligence and biotechnology infrastructure. Both have signed agreements with major Western and Chinese technology firms to deploy advanced systems. Neither has produced specific neuro-rights legislation, but both are likely to encounter the question soon. The form their answers take will be consequential not only for their own citizens but, given the regional influence of those states, for the broader Arab region.

The UNESCO Recommendation on the Ethics of Neurotechnology, adopted in 2025 with support from a majority of Arab states, provides a soft-law anchor that Arab jurisdictions can use to develop their own frameworks. The Recommendation's emphasis on human dignity, free and informed consent, and protection of vulnerable groups translates well across the cultural and religious traditions of the region. Civil-society organizations seeking to advance neuro-rights in Arab jurisdictions are likely to find the UNESCO instrument more productive as a reference than the more politically loaded UN Human Rights Council instruments.

South and Southeast Asia

India, the world's largest democracy and home to more than a billion people, has a constitutional right to privacy recognized by the Supreme Court in the 2017 Puttaswamy judgment. This judgment provides a powerful foundation on which neuro-rights protections could be built. India also has an active neurotechnology research community, a growing biotech industry, and a tradition of public-interest litigation that has, on multiple occasions, brought new technologies under judicial scrutiny.

Despite these favourable preconditions, India has produced limited neuro-rights specific legislation or case law. The Digital Personal Data Protection Act of 2023 provides a general data-protection framework but does not specifically address neural data. Public discussion of neurotechnology in Indian media has been intermittent and is heavily focused on medical applications. As consumer neurotechnology spreads in the Indian market — driven by both domestic startups and foreign products — the absence of specific protections will become increasingly visible.

Southeast Asia presents a more diverse picture. Singapore has a sophisticated technology-policy infrastructure and has begun to engage with the ethics of AI in ways that have implications for neurotechnology. The Philippines has a vigorous civil-society sector with experience advocating for digital rights. Indonesia, Malaysia, Thailand, and Vietnam are each developing data-protection regimes at varying paces. None has yet produced specific neuro-rights legislation, but the regional digital-rights movement is active and increasingly engaged with neurotechnology questions.

The Pacific and Small Island Developing States

Small island developing states are often overlooked in discussions of advanced technology policy. Yet many of these states have produced disproportionate leadership on emerging human-rights questions, particularly through their work in the UN system. Pacific states have been prominent voices in the development of climate-related rights, in the protection of indigenous knowledge, and in the regulation of submarine cables and digital infrastructure. The same diplomatic capacity could be mobilized for neuro-rights.

Several Pacific states have indigenous legal and philosophical traditions that bear directly on neuro-rights questions. The concept of personhood in many Pacific traditions extends beyond the individual to encompass relationships with family, community, ancestors, and land. The legal protection of cognitive sovereignty in these traditions is, in some respects, more developed than in Western legal cultures, even if the specific terminology of "neuro-rights" is not used. Drawing on these traditions could enrich the global conversation in ways that the dominant European and Latin American framings have not.

Why the Asymmetry Matters

Three consequences follow from the current asymmetric distribution of neuro-rights activity. First, the Global South is exposed to neurotechnology deployment without the protective frameworks that more vocal jurisdictions have begun to construct. Consumer EEG headsets sold in Lagos or Nairobi or Manila operate under data-protection regimes that are weaker, or sometimes absent, compared to those in Berlin or São Paulo. The companies producing the products are typically the same. The protections are not.

Second, the absence of Global South voices in the international neuro-rights conversation produces frameworks that imperfectly reflect the experiences and values of the majority of the world's population. The right to mental privacy as articulated in European and Latin American documents is grounded in an individualist conception of personhood that does not always translate cleanly to communal and relational conceptions of selfhood found elsewhere. A truly global neuro-rights framework would integrate these perspectives. The current frameworks do not.

Third, the diplomatic infrastructure for a global treaty depends on broad participation. The Biological Weapons Convention and the Chemical Weapons Convention each succeeded because they recruited support from across the world, not only from a coalition of advanced economies. A neurotechnology treaty that aspires to comparable reach must invest now in the regional and national capacity that will make broad participation possible later. This

means funding research and civil-society work in the Global South, supporting translations and regional convenings, and ensuring that draft treaty texts reflect Global South input from early stages, not as a final consultation.

The neuro-rights movement has, so far, succeeded most where it has had the advantage of intellectual leadership and political will. Extending it to the Global South requires the deliberate construction of those conditions in jurisdictions where they do not yet exist. That is not the work of years; it is the work of the next decade. Whether the global movement succeeds in this construction will determine, more than any other single factor, whether the world arrives in 2035 with a genuinely global framework or with a fragmented one in which the protections of citizens depend on the accident of where they live.

CHAPTER 15

Colorado and California — US State-Level Neurodata Protection

While international bodies debated frameworks and treaties, a parallel movement was gaining momentum within the United States—not at the federal level, but in the statehouses of Colorado and California. In 2024, both states enacted legislation that amended existing privacy laws to explicitly protect neural data as sensitive personal information, marking the first time any US jurisdiction had taken such targeted action.

The Colorado Privacy Act Amendment

Colorado's approach built upon the Colorado Privacy Act (CPA), which had already established a framework for consumer data protection. The 2024 amendment added "neural data" to the CPA's definition of "sensitive data," thereby triggering the Act's enhanced protections. Under the amended law, companies collecting neural data must obtain explicit consumer consent before processing such information, and they must provide clear disclosures about how the data will be used.

The amendment defined neural data broadly to encompass "information that is generated by the measurement of the activity of an individual's central or peripheral nervous system and that can be processed by or with the assistance of a device." This definition was carefully crafted to cover both current technologies—such as EEG headsets and fMRI scans—and future devices that might emerge as the industry evolved.

Critically, the Colorado amendment recognized what the UN report had identified as a core challenge: that neurodata "provide highly sensitive information and allow aspects of individuals' identity and personality to be revealed." The report had warned that such data could expose "sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions." Colorado's legislators took this warning seriously, treating neural data with the same level of protection already afforded to biometric data, health information, and other sensitive categories.

California's AB 1212 and the CCPA

California, which had long positioned itself as a leader in digital privacy through the California Consumer Privacy Act (CCPA) and its successor, the California Privacy Rights Act (CPRA), followed a similar but distinct path. Assembly Bill 1212, signed into law in 2024, amended the CCPA to include "neural data" within the definition of "biometric information."

The California approach had particular significance because the CCPA, unlike most state privacy laws, included a private right of action for certain violations. This meant that Californians whose neural data was collected or used without proper consent could sue companies directly, rather than relying solely on enforcement by the state attorney general. The amendment also required businesses to provide specific notice before collecting neural

data and to obtain affirmative authorization for any use beyond what was reasonably necessary to provide the requested service.

California's action reflected the urgency that the UN report had identified: "The risk is that, without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations." By embedding neural data protection into an existing enforcement mechanism, California created a model that other states could adapt.

The Significance of State-Level Action

The Colorado and California amendments were notable not only for their content but for what they represented. In the absence of federal legislation—the United States had no comprehensive national privacy law that could address neurotechnology—states were forced to act on their own. The result was a patchwork of protections, but one that established important precedents.

Both states recognized that existing privacy frameworks, designed for an era of data collection from computers and smartphones, were inadequate for the unique challenges posed by neurotechnology. As the UN report had noted, "The specificities of brain data may not be adequately protected by existing international and domestic privacy protection frameworks." The report called for "context-specific standards to protect against the non-consensual intrusion of third parties."

Colorado and California answered that call within their jurisdictions. Their actions demonstrated that neurodata protection was not merely an abstract human rights concern but a practical regulatory challenge that could be addressed through existing legal mechanisms—provided legislators had the will to act.

Gaps and Limitations

Despite their significance, the state-level amendments had important limitations. Neither law addressed the most invasive forms of neurotechnology, such as implanted devices that could write to the brain as well as read from it. The protections focused on data collection and processing, not on the direct neural interventions that the UN report had flagged as potentially violating "freedom of thought" and "personal integrity."

Moreover, the state laws could not address the national security and military applications of neurotechnology that the report had identified as particularly concerning. The UN document had warned of "coercive uses of neurotechnologies to augment resilience on the battlefield" and the risk that "neurotechnologies can be used to interfere and manipulate individuals" in ways "similar to 'brainwashing.'" State privacy laws, no matter how robust, had no jurisdiction over federal agencies or military programs.

Nevertheless, the Colorado and California amendments represented a crucial first step. They established that neural data was different from other forms of personal information—more intimate, more revealing, and deserving of heightened protection. As other states considered similar legislation, and as advocates pushed for federal action, these state-level models

provided both a template and a proof of concept: neurodata protection was legally feasible, politically achievable, and urgently necessary.

CHAPTER 16

The EU AI Act and Its Gaps Regarding Neurotechnology

The European Union's Artificial Intelligence Act, adopted in 2024, represents the world's first comprehensive horizontal regulation of artificial intelligence. It classifies AI systems according to risk, with "unacceptable risk" systems prohibited entirely, "high-risk" systems subject to strict obligations, and "limited" and "minimal" risk systems subject to lighter requirements. Among the prohibited practices are the deployment of "subliminal, manipulative, or deceptive techniques to distort behavior and hinder an individual's informed decision-making," as well as social scoring systems and remote biometric identification in certain contexts.

These provisions might appear, at first glance, to offer meaningful protection against the most dangerous applications of neurotechnology. A brain-computer interface that manipulates a user's emotions without their knowledge, for example, would seem to fall squarely within the prohibition on subliminal manipulation. An AI system that uses neural data to make employment or insurance decisions could be classified as high-risk, triggering requirements for risk assessment, transparency, and human oversight.

Yet the AI Act contains a fundamental flaw that renders it inadequate as a framework for protecting neuro-rights: it says nothing about neurotechnologies that are not AI-driven. Implants, directed energy weapons, and non-invasive mind-reading devices that do not rely on machine learning algorithms fall entirely outside the Act's scope. A consumer EEG headband that records brain activity and transmits the raw signal to a cloud server—without any AI processing—would not be regulated as a high-risk AI system, even though it collects the most intimate data a person can generate.

The human-rights advocacy contribution to the UN Office of the High Commissioner for Human Rights, submitted in March 2026, makes this point explicitly: the AI Act "says nothing about other neurotechnologies such as implants, directed energy weapons (DEW) or brain-computer interfaces (BCI)." This is not a minor oversight. It is a structural gap that leaves the most invasive forms of neurotechnology—those that can read thoughts remotely, write data back into the brain, or cause neurological harm without the target's awareness—entirely unregulated at the European level.

Consider the case of directed energy weapons. The AI Act, focused as it is on software systems, has no provisions addressing the use of electromagnetic radiation or ultrasound to affect neural function. Yet the Havana Act of 2021, passed by the U.S. Congress, acknowledged that such weapons had caused "unusual cognitive and neurological disorders" in diplomats, military personnel, and CIA agents serving abroad since 2016. The Active Denial System, developed by the U.S. armed forces and private companies, is a directed energy weapon that uses millimeter waves to cause intense pain. The AI Act provides no framework for assessing, prohibiting, or even acknowledging such technologies.

Similarly, non-invasive brain-computer interfaces that operate at a distance—sensors that can detect neural activity without direct contact with the user's body—fall outside the Act's regulatory scope. The European Data Protection Supervisor has acknowledged that "there are neural sensors that are not in direct contact with the user's body, but are located remotely, so the user may not be aware of their presence." The AI Act does not address this category of devices at all, leaving a regulatory vacuum precisely where the risk of covert surveillance is highest.

The Act's focus on AI also means that it fails to address the most fundamental question raised by neurotechnology: not whether an algorithm makes decisions based on neural data, but whether the data should be collected in the first place. The AI Act regulates what AI systems do with data; it does not regulate the act of reading the brain itself. A device that records neural activity without any AI processing is simply a data collection tool, and the AI Act has nothing to say about it.

This gap is particularly significant because the European Union has no dedicated neuro-rights instrument. It relies instead on a patchwork of the GDPR, the AI Act, and a 2019 OECD recommendation, "none of which was written with brain-reading consumer devices in mind," as TI One Voice's Stephane van der Aa noted in his November 2025 letter to members of the European Parliament. The GDPR, adopted in 2018, predates the commercial proliferation of consumer neurotechnology and does not contain a specific category for "neural data." While neural data can be classified as sensitive health or biometric data under the GDPR, this classification is imperfect: it requires the data to be processed for a specific purpose, and it does not address the unique risks posed by devices that can both read and write to the brain.

The Council of Europe Framework Convention on Artificial Intelligence, also adopted in 2024, represents the first binding international treaty on AI. It mandates risk and impact assessments and provides safeguards such as the right to challenge AI-made decisions. But it suffers from the same limitation as the EU AI Act: it applies only to AI systems, not to neurotechnology more broadly. Moreover, the Convention excludes activities related to national security and defense—precisely the domain where the most concerning neurotechnology applications are being developed.

Chile's approach, by contrast, demonstrates what a comprehensive regulatory framework might look like. The 2021 constitutional reform protects "brain identity as well as the information derived from it" and gives brain data the same legal status as a human organ, prohibiting its sale, trafficking, or manipulation. Unlike the EU AI Act, Chilean law "does not focus on AI as the only neurotechnology to be regulated, and also legislates on devices that can read thoughts and emotions remotely, and that know how to implement data in return."

The gap in the EU framework is not merely theoretical. As neurotechnology moves from medical settings into consumer markets, workplace monitoring, education, and security, the absence of specific protections for neural data and neural interventions creates a regulatory vacuum that private actors—and potentially state actors—can exploit. Devices for non-medical use, including those used for work, education, or leisure, are "often not subject to

oversight with regard to their effectiveness, safety or ethical implications," as the human-rights advocacy submission notes.

The AI Act is a significant achievement in the regulation of artificial intelligence. But it is not, and was never intended to be, a comprehensive framework for neuro-rights. The European Union now faces a choice: either amend the AI Act to cover all neurotechnologies, regardless of whether they employ AI, or develop a separate, dedicated instrument that addresses the full spectrum of neurotechnology risks—from consumer EEG headbands to directed energy weapons, from non-invasive mind-reading to implanted devices that can write to the brain. Without such action, the most intimate and vulnerable domain of human experience—the inner realm of thought—remains unprotected by European law.

CHAPTER 17

The Council of Europe Framework Convention on AI

In September 2024, the Council of Europe adopted the Framework Convention on Artificial Intelligence—the first-ever binding international treaty in this domain. Signed by the European Union, the United States, Japan, Canada, and approximately twenty other nations, the Convention represents a significant step toward global governance of AI. Yet, like the EU AI Act, its relevance to neuro-rights is circumscribed by its scope and exclusions.

The Framework Convention mandates risk and impact assessments to mitigate potential harms from AI systems and provides procedural safeguards, including the right to challenge decisions made by AI. Implementation is overseen by a Conference of the Parties. These mechanisms are valuable for AI governance generally, but they do not address the specific challenges posed by neurotechnology.

The Convention's most significant limitation for neuro-rights is its focus on artificial intelligence as the object of regulation. As the human-rights advocacy submission to the UN High Commissioner for Human Rights observes, the Council of Europe framework—like the EU AI Act—"does not focus on AI as the only neurotechnology to be regulated." It does not cover devices that can read thoughts and emotions remotely, nor does it address implants, directed energy weapons, or brain-computer interfaces that do not rely on AI in their core functioning.

The exclusion of national security and defense activities further limits the Convention's reach. As with the EU AI Act, the most concerning neurotechnology applications—military brain-computer interfaces, cognitive warfare capabilities, directed energy weapons—fall outside the treaty's scope entirely.

For the Convention to become law, it must be ratified by at least five signatory governments. As of early 2026, this threshold had not yet been reached, meaning the treaty does not yet have the force of law. Even when ratified, its provisions will apply primarily to public authorities and private entities acting on their behalf, leaving consumer neurotechnology products largely unregulated.

The Council of Europe Framework Convention is a meaningful achievement for AI governance. But it is not a neuro-rights instrument. It does not protect mental privacy, cognitive liberty, or personal identity from the full spectrum of neurotechnology threats. The gap it leaves is one that dedicated neuro-rights legislation—at the national, regional, and international levels—must fill.

CHAPTER 18

Comparative Analysis — A Patchwork of Protections

After six chapters of jurisdiction-by-jurisdiction examination, it is useful to step back and ask what, in aggregate, the world has produced. The answer is what the title of this chapter calls a patchwork: a heterogeneous collection of instruments, drafted at different times, in different legal traditions, with different ambitions and different gaps. This chapter compares the leading instruments side by side, identifies the patterns and the discrepancies, and assesses what additional pieces the patchwork still requires before it can be called a quilt.

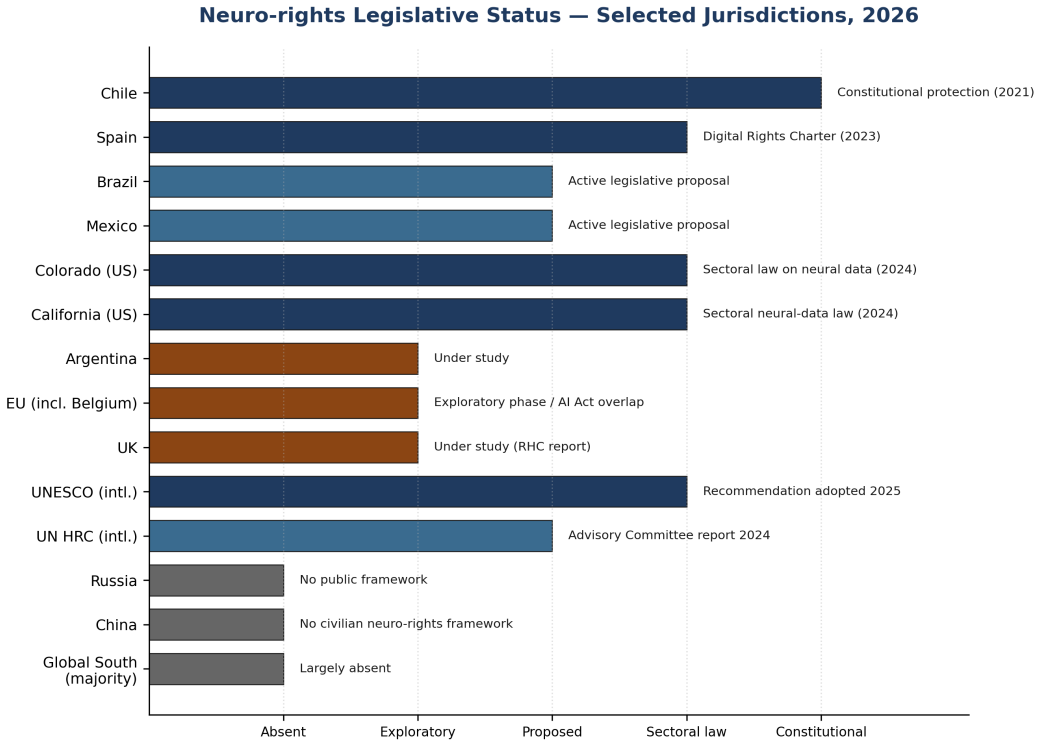


Figure 12.1. *Neuro-rights legislative status across selected jurisdictions in 2026. The visible patchwork reveals both leadership clusters (Latin America, Iberia, US sectoral law) and large gaps (Global South, authoritarian states, multilateral coverage).*

Dimensions of Comparison

Useful comparison requires consistent dimensions. The instruments surveyed in this section vary along several axes: their legal form (constitutional, statutory, regulatory, soft-law), their scope (general data protection, specific neural-data provisions, or full neuro-rights frameworks), their definition of neural data (raw signals only, or inferences as well), their treatment of consent (general, heightened, or absolute), their enforcement mechanisms (administrative, civil, criminal), and their geographic reach (national, regional, or international). A coherent assessment of the global state of protection requires evaluating each instrument along each axis.

Instrument	Legal Form	Defines Neural Data	Inferences Covered	Enforcement	Geographic Scope
Chile constitutional amendment	Constitutional	Yes (general)	Implicit	Constitutional review	National
Spain Charter of Digital Rights	Soft-law executive	Yes (general)	Implicit	Interpretive only	National
Brazil PL 4.690 (pending)	Constitutional + statutory	Yes (proposed)	Yes (proposed)	Administrative + civil	National
Mexico Art. 4 proposal	Constitutional (pending)	Partial	Partial	Constitutional review	National
Colorado HB24-1058	Statutory amendment	Yes (raw signals)	No	AG enforcement	State-level
California CCPA amendment	Statutory amendment	Yes (raw signals)	Limited	AG + private right	State-level
EU AI Act	Regulation	Indirect (via 'biometric')	Partially	Administrative	EU-wide
CoE Framework Convention on AI	Treaty (general)	No	No	State implementation	European/Open
UN HRC Resolution 51/3 + Report	Soft-law	Conceptual	Conceptual	None directly	Global (advisory)
UNESCO 2025 Recommendation	Soft-law	Yes	Yes	State implementation	Global (advisory)

Definitional Inconsistency

The most consequential inconsistency among existing instruments is the definition of neural data itself. The U.S. state laws (Colorado, California) define neural data narrowly, focusing on raw electrical signals collected by neurotechnology devices. The EU AI Act covers biometric and emotion-recognition data and reaches neural data indirectly, through that more general category. The Latin American instruments tend to define neural data more broadly, including derived inferences. The UNESCO Recommendation follows the broader Latin American framing. The result is that what counts as protected neural data depends

entirely on the jurisdiction. A user generating signals from a consumer headset in Brazil is, in principle, protected against inferences drawn from those signals; the same user generating the same signals in Colorado is not.

This inconsistency is exploitable. A neurotechnology company that operates internationally can choose to process raw signals in a jurisdiction with broad protections and derive inferences in a jurisdiction with narrow ones. The signals themselves never leave the protected jurisdiction. The inferences — the more sensitive and more usable product — are produced in jurisdictions where the law does not reach. The legislative recommendation that follows is straightforward: every neuro-rights instrument should define neural data to include the algorithmic inferences drawn from neural signals, and should make clear that the inferences are protected regardless of where in the world they are produced.

Consent Regimes

A similar inconsistency afflicts the treatment of consent. Some instruments treat neural-data processing as requiring ordinary consent under the relevant data-protection regime. Others require heightened consent, with explicit acknowledgement of the special nature of the data. A few — most notably the Chilean amendment, in spirit if not in operational detail — suggest that certain uses are absolutely prohibited regardless of consent. The variation matters because consent is the principal lever by which neurotechnology companies legitimize data collection, and weakening consent requirements is the principal mechanism by which protection is eroded over time.

A reasonable target standard for consent in the neurotechnology context has four elements: it must be explicit (not inferred from the use of a product), informed (with disclosures sufficient for genuine understanding), specific (limited to articulated purposes), and revocable (with effective procedures to withdraw and to require deletion). Many existing instruments achieve some of these elements; few achieve all four; and several do not articulate the requirements clearly enough to be operationally meaningful. The trend, observable across multiple jurisdictions, is toward strengthening consent. The pace of strengthening is, however, slower than the pace of technological development.

Enforcement and the Verifiability Problem

All neuro-rights instruments confront the same enforcement challenge: it is difficult to prove that neural data has been improperly collected or used. The technical signature of a violation is faint. The data, once collected, can be processed in jurisdictions where local enforcement does not reach. Inferences derived from neural data are difficult to distinguish from inferences derived from other sources. Even where a violation is detected, the chain of custody required to establish it in court is fragile.

These challenges argue for several enforcement innovations that few current instruments adopt. First, mandatory logging by neurotechnology devices, in tamper-evident form, of every data transmission and every inference operation. Second, third-party audit rights for regulators and, in appropriate circumstances, individual subjects. Third, presumptions of unlawful processing where logs are absent or incomplete. Fourth, criminal as well as civil

sanctions for the most serious violations, particularly those involving non-consensual modulation of cognitive function. Of these innovations, only the third is implicit in some existing data-protection regimes; the others are largely absent from current law. Adding them is a key task for the next generation of drafting.

The Coverage Gap: Inferences, States, and Weapons

When the existing instruments are laid side by side, three coverage gaps become visible. The first is the inferences gap, addressed above: most instruments protect raw neural data but few clearly protect inferences. The second is the state-actor gap: many instruments focus on consumer products and commercial data processing, but few clearly constrain state agencies acting in security, law-enforcement, or military capacities. The third is the weapons gap: virtually no existing instrument explicitly prohibits the use of directed-energy and other neurotechnological systems as weapons against persons, though several international instruments could be read to support such a prohibition.

Closing these gaps requires a different generation of instruments than the ones currently in force. Specifically, it requires explicit treatment of state actors, explicit prohibitions on weaponization, and explicit inclusion of inferential data within the definition of neural data. These elements together do not require a new global treaty, though such a treaty would help. They could be achieved through revisions of existing national and regional instruments, through new draft model legislation that civil-society organizations and academic centres can propagate, and through interpretive guidance from oversight bodies that exists already.

Toward Convergence

The pattern of recent activity suggests that convergence is possible. Successive instruments tend to incorporate language and definitions from their predecessors, with refinements based on observed gaps. The Latin American instruments built on each other. The European instruments increasingly draw on the Latin American ones. The UNESCO Recommendation drew on both, plus the academic literature, plus consultations with states and civil society from across the world. A draft model law produced today and circulated through the academic and civil-society networks that connect these jurisdictions can plausibly influence the next round of national legislation everywhere from Madrid to Mexico City to Kuala Lumpur.

The patchwork is not yet a quilt. But the threads are being drawn together. Whether the final fabric is dense and protective, or loose and porous, depends on the choices made in the next several years by the drafters of the next round of instruments. Those choices are not technical; they are political. The most consequential of them — the inclusion of inferences, the constraint of state actors, the prohibition of weaponization — are the choices the existing patchwork most clearly identifies as outstanding. Filling them is the central task this guide returns to in the chapters that follow.

CHAPTER 18-bis

Existing Laws at Every Level — A Consolidated Reference

The previous chapters of Part II surveyed the international and selected national instruments that already address neurotechnology. This chapter consolidates that material into a single, level-by-level reference. It is intended for the legislator, civil servant, lawyer or citizen who needs to know, in one place, what protection a given act of neural reading or writing already attracts under the law as it stands in 2026, and where the protection runs out.

The exercise is deliberately defensive. We do not assume that what exists is sufficient — Parts XII and XII-bis will argue at length that it is not — but we do insist that the case for new instruments be built on an honest map of the old. A claim that "no law exists" is now almost always wrong; a claim that "existing law is sufficient" is reliably wrong in a different direction. The truth is more interesting: a partial, overlapping, jurisdictionally uneven coverage that catches some acts of neural intrusion squarely, others sideways through general-purpose instruments, and others not at all.

Figure 18bis.1 visualises the picture across twelve protective dimensions and fourteen representative jurisdictions. The reader will see immediately that the brightest spots — strong, dedicated protection — cluster around raw neural data in a handful of pioneer jurisdictions, and that whole columns and rows remain dark: neuro-weapons, black-box logging, and most of the Global South.

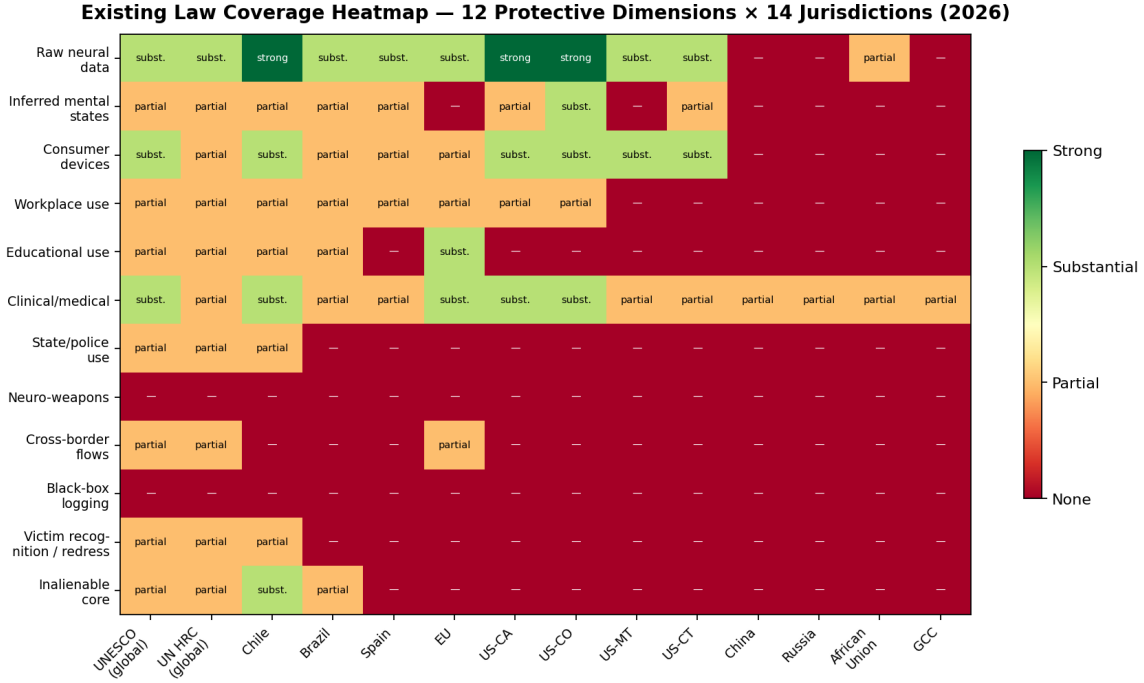


Figure 18bis.1. Existing law coverage heatmap. Twelve protective dimensions across fourteen representative jurisdictions in 2026. Red indicates no dedicated protection; green indicates substantial or strong protection. The visible pattern is a thin band of strong protection across the consumer-data dimensions in Chile, California, and Colorado, partial protection across most jurisdictions for clinical and medical uses, and an almost-complete absence of protection for neuro-weapons, black-box logging, and Global-South jurisdictions.

The Six Tiers in Practice

Existing protection of the mind is delivered through six legally distinct tiers, each with its own pathologies of over- and under-protection. The honest summary is that none of them was written with brain-reading consumer devices in mind, and the fit is awkward at every level.

Tier 1 — Supra-national Soft Law and Treaty Norms

At the apex sits a small but expanding body of supra-national instruments. The 2025 UNESCO Recommendation on the Ethics of Neurotechnology (covered in Chapter 11) is the first global standard dedicated to the field; the 2022 UN Human Rights Council Resolution 51/3 (Chapter 9) and the 2024 UN A/HRC/57/61 report (Chapter 10) have given the international human-rights system a concrete neurotechnology mandate. The 2019 OECD Recommendation on Responsible Innovation in Neurotechnology was the first international instrument of any kind, and although it is non-binding it has shaped national policy in member states. The Council of Europe Framework Convention on AI (Chapter 17) covers neural inferences indirectly via its AI provisions.

Above and behind these neurotech-specific instruments stands the general international human-rights canon. Article 18 of the International Covenant on Civil and Political Rights protects freedom of thought, conscience and religion as an absolute, non-derogable right. Article 17 protects privacy. Article 7 prohibits torture and cruel, inhuman or degrading treatment. The Convention against Torture, the Convention on the Rights of the Child and the Convention on the Rights of Persons with Disabilities all bear on neurotechnology even though none of them mentions it by name.

Outside human-rights law, three weapons-related instruments are increasingly relevant. The 1972 Biological Weapons Convention and the 1993 Chemical Weapons Convention together established the principle that certain categories of capability are inadmissible regardless of military utility, and have been suggested as models for a future neuroweapons convention (see Chapter 38). The 1980 Convention on Certain Conventional Weapons, with its protocols on blinding lasers and incendiaries, shows that the international community can prohibit specific weapon types ex ante on the basis of disproportionate or treacherous harm.

Trade law completes the picture. The WTO Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS) governs neuro-patent flows; the GATS framework regulates cross-border services that may include neurodata processing; and a still-emerging body of digital-trade rules in plurilateral agreements such as the CPTPP and the Digital Economy Partnership Agreement will determine how neural data crosses borders. None of these were drafted with the brain in mind. All will need to be interpreted, or amended, to account for it.

Tier 2 — Regional Frameworks

The regional tier is the most uneven. The European Union, despite the absence of a dedicated neuro-rights instrument, already supplies several layers of protection through the General Data Protection Regulation (which treats biometric and health data as a special category requiring explicit consent for processing), the AI Act of 2024 (which prohibits certain emotion-recognition uses in workplaces and schools and bans specific manipulative subliminal techniques — see Chapter 16), the Medical Device Regulation 2017/745, the Digital Services Act, and the Network and Information Security Directive 2 (NIS2). The Charter of Fundamental Rights of the European Union, in Articles 1, 3, 7, 8 and 10, supplies the constitutional-level backstop.

The Council of Europe complements this with the European Convention on Human Rights (Articles 8, 9 and 14 are the most relevant), the Convention 108+ on data protection, the Oviedo Convention on biomedicine, and the new Framework Convention on AI (Chapter 17). The European Court of Human Rights has not yet ruled on a dedicated neurotechnology matter, but its jurisprudence on bodily integrity (Pretty v. UK), medical autonomy (V.C. v. Slovakia), and the absolute prohibition on torture (Selmouni v. France) provides a doctrinal scaffold.

Latin America has the world's most rights-progressive regional posture. The Organization of American States and the Inter-American Commission on Human Rights have produced declarations on neurotechnology, and the Inter-American Court of Human Rights' jurisprudence on the right to personal identity (Atala Riffo) and to mental integrity (Ximenes Lopes) is unusually well-suited to neurotech adaptation. The pioneering work of Chile, Brazil, Spain, Mexico, Colombia, Argentina and Uruguay (Chapter 12) makes the region the world's leading norm entrepreneur.

The African Union has produced the Malabo Convention on Cyber-Security and Personal Data Protection (in force 2023), which covers some neurodata-processing activities but does not address neural inferences or workplace surveillance. The 1981 African Charter on Human and Peoples' Rights (Articles 4, 5 and 6) supplies the human-rights ceiling. The Arab Charter on Human Rights (2004) covers privacy and bodily integrity, though enforcement is uneven. ASEAN has a Framework on Personal Data Protection (2016), revised in 2024 to include biometric data, but no neurotech-specific instrument.

Tier 3 — National Constitutional Protection

At the constitutional tier, three jurisdictions stand alone in having explicit textual recognition of neuro-rights. Chile's Law 21.383 of 2021 amended Article 19 of the Constitution to require that scientific and technological development respect the psychic integrity of persons; the Constitutional Court has since interpreted this in Emotiv (Chapter 8) to support a remedial order requiring data deletion. Brazil's 2024 constitutional amendment on mental integrity follows the same template. Spain's 2021 Charter of Digital Rights, while not constitutional in the formal sense, has constitutional-adjacent status as guidance for legislators interpreting Articles 18 (privacy) and 15 (physical and moral integrity) of the 1978 Constitution.

Beyond these textual recognitions, every constitutional democracy supplies some level of indirect protection. The First Amendment of the United States Constitution protects freedom of thought and conscience; the Fourth, Fifth and Fourteenth Amendments protect against unreasonable search and self-incrimination, with strong implications for compelled brain reading. The German Basic Law's Article 1 (human dignity) and Article 2 (free development of personality) provide unusually strong dignity-based protection. The French Constitution and the 1789 Declaration of the Rights of Man supply liberty-of-thought protection. The Indian Constitution's Article 21 has been read by the Supreme Court (Puttaswamy, 2017) to include a robust right to privacy that extends to mental autonomy.

What no constitution yet does — including Chile's — is treat neural data and inferences as an inalienable category that cannot be waived even with consent. The closest analogue is Germany's doctrine of human dignity (Menschenwürde) under Article 1, which the Federal Constitutional Court has treated as non-derogable. Whether that doctrine will be extended to neural privacy is one of the open questions of the next decade.

Tier 4 — National Statutory Protection

The national statutory tier is where most of the action is happening. The European Union's GDPR, the United Kingdom's Data Protection Act 2018, Brazil's LGPD, Japan's Act on the Protection of Personal Information, South Korea's PIPA, South Africa's POPIA, Singapore's PDPA and Canada's PIPEDA all treat health and biometric data as a heightened category. None of them was drafted with neural data specifically in view; all of them can be read to cover at least raw neural signals; most of them leave inferences in a gray zone.

Medical-device statutes supply a second layer. The US Food, Drug, and Cosmetic Act and FDA's 21 CFR 882 (Neurological Devices) govern medical neurotechnology; the EU Medical Device Regulation does the same in Europe. Both regimes require pre-market review of high-risk devices but do not require black-box logging of what the device reads from or writes to the brain. Both also leave consumer wellness wearables almost entirely unregulated, because such devices are deliberately marketed as non-medical to escape this layer.

Workplace law supplies a third. The US Genetic Information Nondiscrimination Act (GINA), the EU Workplace Health and Safety Directive, the various national equality acts, and ILO

Conventions 158 and 187 each prohibit discrimination on some bases that neurodata might create or reveal. The AI Act's Article 5 specifically bans workplace emotion-recognition AI systems in the EU. But none of these statutes addresses the more subtle question of attention monitoring or cognitive-load measurement directly.
Criminal procedure law supplies a fourth. The Fifth Amendment privilege against self-incrimination in the US, Article 6 of the European Convention on Human Rights in Europe, and analogous provisions worldwide arguably prohibit compelled brain reading, but no court has yet ruled definitively on this point. Wisconsin v. Cucalon and a small number of analogous lower-court rulings hint at the direction; conclusive doctrine is still missing.

Tier 5 — Sectoral Regulation

Below the statute, every regulated sector adds its own layer. Medical-device regulators issue product-specific approvals. Telecommunications regulators license radio-frequency emissions and could in principle attribute illegal devices. Education ministries control what may be measured about children in classrooms. Defence ministries govern the military uses described in Part IV. Insurance regulators determine whether neurodata can be used in underwriting (most explicitly prohibit it for now). Financial regulators set the limits on neural-marketing for credit products.
The sectoral tier is where compliance actually happens. A constitutional clause or a national statute is too abstract to bind a product manager; only the granular rule — what label must appear on a packaging, what test must be passed, what log must be kept — actually constrains conduct in the field. This is also where the regulatory holes are deepest, because there is rarely a regulator with the technical staff to write neurotech-specific rules in any single sector.

Tier 6 — Sub-national and Municipal

Finally, at the local level, US states have produced the most innovative dedicated neuro-legislation of the past two years. Colorado's 2024 amendment to its Privacy Act classifies neural data as sensitive data, requires explicit consent for collection, and creates a private right of action. California's 2024 amendment to the CCPA does the same. Montana's 2025 Genetic Information Privacy Act amendment extends coverage to neural data. Connecticut's 2025 Data Privacy Act amendment does the same. Approximately fifteen further US state bills are pending.
Sub-national protection extends beyond the US. Quebec's Law 25 imposes some of the world's strongest data-protection rules. The German Länder have their own Landesdatenschutzgesetze. Brazilian states have their own consumer-protection regimes. Some municipalities have begun banning specific deployments: the San Francisco municipal ban on city-government use of facial recognition (a precursor to neural surveillance) has been extended in Berkeley and Oakland.

How the Layers Interact

The honest assessment of the layered system is that it works best when all six tiers reinforce each other and worst when any single tier is treated as sufficient on its own. A national statute is only as strong as the regulator empowered to enforce it; a regulator is only as effective as the sectoral rules that translate the statute into product specifications; sectoral rules are only as protective as the constitutional ceiling above them and the international floor beneath them; and the international floor is only as binding as the verification regime that gives it teeth.
The Chilean Emotiv case (Chapter 8) is the only example to date of all the layers working together. Chile had the constitutional ceiling. The Constitutional Court was empowered to interpret it. A constitutional clause aimed at psychic integrity was matched to ordinary data-protection law to deliver a concrete remedy: deletion of the affected data. International norms (the UNESCO Recommendation, the OECD framework, the UN Human Rights Council resolution) supplied the interpretive backdrop. The case stands as a proof of

concept: layered protection can work in practice. It also stands as a warning: the layers worked because Chile had built them. Most jurisdictions have not.

The Coverage Gaps

Reading across all six tiers, six structural coverage gaps stand out. They are the principal targets of the ideal-legislation chapter (Chapter 67-bis) later in this book.

Gap one: inferred mental states. Most regimes protect the neural signal and ignore the algorithmic inference. California's amendment explicitly excludes data derived from neural signals by an algorithm; Montana's excludes "downstream physical effects"; Colorado is one of the few that explicitly includes inferences. A law that protects the lock but not the valuables is not a law that protects the valuables.

Gap two: 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. The EU AI Act's narrow ban on workplace and school emotion-recognition is one of the few prominent exceptions; even it is riddled with exemptions.

Gap three: manipulation and output. 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. The EU AI Act's prohibition of certain subliminal manipulation techniques is the only consequential rule on writing in the entire global stack.

Gap four: 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. The Chilean Emotiv case produced one remedy (data deletion); no jurisdiction has yet awarded damages for unauthorised neural reading.

Gap five: detection and audit standards. No jurisdiction yet requires neurotech systems to keep tamper-evident logs of what they read and wrote. Without such logs, abuse is unprovable after the fact even when it has clearly occurred.

Gap six: weaponisation. No binding international instrument prohibits the development, production, stockpiling, transfer or use of neuroweapons. The European Parliament's 1999 resolution called for such a ban; nothing has come of it. This is the largest hole of all, and the principal subject of Chapter 38 and the proposed convention drafted in Chapter 68.

A Jurisdictional Quick Reference

The following table is intended as a quick-reference for the practitioner who needs to know what regime applies in a given jurisdiction. It is necessarily simplified; the reader should consult the underlying instruments before relying on this summary for any consequential decision.

Jurisdiction	Constitutional	Statutory	Sectoral	Sub-national
Chile	Art. 19 (Law 21.383)	Personal Data Protection Law	Health, education	Municipal consumer rules
Brazil	Mental-integrity amendment	LGPD	ANVISA, education	State laws (SP, RJ)
Spain	Art. 18 + Charter of Digital Rights	LOPDGDD	AEMPS, workplace	Catalan, Basque rules
EU member states	EU Charter Arts. 1,3,7,8,10	GDPR, AI Act, MDR	National regulators	Sub-national variants
UK	Human Rights Act 1998	DPA 2018	MHRA, HSE	Devolved nations
USA federal	I, IV, V, XIV Amendments	FTC Act, MIND Act (pending)	FDA, FCC, EEOC	— (state level dominant)
US states (CA, CO, MT, CT)	—	State privacy acts	State regulators	Municipal
Canada	Charter ss. 7, 8	PIPEDA,	Health Canada	Provincial (QC

		provincial acts		Law 25)
Germany	Basic Law Arts. 1, 2	BDSG, IfSG	BfArM, BNetzA	Länder laws
Japan	Constitution Art. 13	APPI	PMDA, MHLW	Prefectural rules
South Korea	Constitution Art. 17	PIPA	MFDS	Provincial
Singapore	Common-law privacy	PDPA	HSA, MOM	—
Australia	Implied rights	Privacy Act 1988	TGA, FWO	State acts
China	—	PIPL, DSL, CSL	NMPA, MIIT	Provincial pilot zones
Russia	—	Federal Law on Personal Data	Roskomnadzor	—
India	Art. 21 (privacy)	DPDP Act 2023	CDSCO	State rules
African Union	African Charter Arts. 4, 5	Malabo Convention	Health ministries	National
ASEAN	National	ASEAN Framework + national	National regulators	National
GCC	National	Patchwork	National regulators	—

A Working Conclusion

The most accurate one-sentence summary of the law as it stands in 2026 is this: every act of neural reading or writing already attracts some legal protection somewhere on the planet, but no jurisdiction yet protects the full set of acts that a determined adversary could perform, and no single jurisdiction has built the layered system that would be required to do so. The international floor is soft. The constitutional ceiling is thin and exists in only three jurisdictions. The statutory middle is the strongest layer but is uneven, definition-fragmented, and largely silent on inferences. The sectoral layer is patchy and short of technical capacity. The sub-national layer is innovative in US states and a few Canadian provinces and almost absent elsewhere.

The implications for the next decade are sharp. Closing the six gaps identified above will require action at all six tiers simultaneously. A national constitutional clause is necessary but insufficient; a sectoral product specification is necessary but insufficient; an international treaty is necessary but insufficient. The argument of Chapters 66 through 68 — and of the new Chapter 67-bis later in this book — is that all of them, together, are required. The remainder of Part II provides the doctrinal foundation. The remainder of the book provides the gap analysis and the drafting.

PART III

The Sovereignty of Thought — Mental Privacy and the Forum Internum

The inner realm of human thought has long been considered inviolable. As the United Nations Human Rights Council Advisory Committee observed in its 2024 report on neurotechnology, "for a long time, human mind and conscience were considered the container of the inner self, a fortress, inaccessible from the outside." This fortress—the *forum internum*—is the psychological space where thoughts, beliefs, and convictions are formed, sheltered from external interference. It is the origin of all that makes us persons: our capacity to reason, to choose, to imagine, to dissent.

The concept of the *forum internum* has deep roots in international human rights law. Freedom of thought, enshrined in Article 18 of the International Covenant on Civil and Political Rights, is designed to protect this inner realm. As the UN report emphasizes, freedom of thought "does not allow for any interference whatsoever and must be protected unconditionally, which explains why it is characterized as an absolute right." Unlike other rights that may be balanced against competing interests, the right to think one's own thoughts, free from impermissible external influence, admits no exceptions.

Neurotechnology fundamentally challenges this ancient protection. For the first time in human history, technologies exist that can penetrate the fortress of the mind. The UN report identifies several characteristics of neurotechnology that make it uniquely disruptive: it can "enable the exposition of cognitive processes," "enable the direct alteration of a person's mental processes and thoughts," "bypass the individual's conscious control or awareness," and "enable non-consensual external access to thoughts, emotions and mental states."

These capabilities strike at the very foundation of human dignity. The *forum internum* is not merely a legal construct—it is the seat of personhood, the space where identity is formed and agency is exercised. When neurotechnology allows external actors to read mental states, manipulate emotions, or alter cognitive processes without consent, it violates the core of what it means to be human.

The UN report warns that such interference can occur without the individual's knowledge. Consumer neurotechnology devices, neuromarketing strategies, and immersive computing systems can all be used to condition the formation of opinions, influence decision-making, and drive compulsive behaviors. The report specifically notes that "neurogaming" devices can drive "compulsive or addictive use, particularly when neural reward circuits target altering consumption behaviors."

The sovereignty of thought is therefore not an abstract philosophical concept—it is a practical necessity for democracy, autonomy, and human flourishing. Protecting the *forum internum* requires recognizing that the mind is the last private place, and that no technology, no commercial interest, and no state security apparatus should be permitted to breach its walls without the most rigorous legal and ethical safeguards.

This section examines the legal foundations of mental privacy, the nature of the *forum internum* under international law, and the specific threats posed by modern neurotechnology. It argues that the absolute character of freedom of thought must be reaffirmed and operationalized in the age of brain-reading and mind-steering technologies.

CHAPTER 19

The Forum Internum — Why Your Thoughts Are Sacred

The concept of the *forum internum*—the inner realm of thought, conscience, and belief—is one of the oldest and most fundamental protections in human rights law. It is the psychological space where individuals develop their thoughts, form their convictions, and exercise their capacity for autonomous decision-making. The United Nations Human Rights Council Advisory Committee, in its landmark 2024 report on neurotechnology, explicitly identified this inner sphere as the primary locus of concern: "Freedom of thought is meant to protect the so-called *forum internum*, that is the inner, psychological realm of the person, where thoughts and convictions are formed."

This protection is not merely aspirational. Under international law, freedom of thought is characterized as an absolute right. The UN report emphasizes that "freedom of thought does not allow for any interference whatsoever and must be protected unconditionally." This absolute character reflects the foundational role that thought plays in human dignity and democratic society. Without the ability to think freely, without fear of external surveillance or manipulation, no other rights can be meaningfully exercised.

The Legal Foundation of Mental Sovereignty

The *forum internum* finds its legal expression in Article 18 of the International Covenant on Civil and Political Rights (ICCPR), which protects freedom of thought, conscience, and religion. The UN Human Rights Committee's General Comment No. 22 (1993) describes this right as "far-reaching and profound," encompassing all matters of personal conviction and belief. The *forum internum* is distinct from the *forum externum*—the manifestation of thought through expression, opinion, or action. While the *forum externum* may be subject to certain limitations (such as those necessary for public order or national security), the *forum internum* itself is inviolable.

The UN report on neurotechnology makes clear that this distinction is critical in the age of brain-reading devices: "Because of the essential role it is called upon to play in democracy, freedom of thought does not allow for any interference whatsoever and must be protected unconditionally." This means that any technology that penetrates the *forum internum*—whether to read thoughts, infer mental states, or manipulate cognitive processes—engages the absolute core of human rights protection.

The Threat of Neurotechnology to the Inner Realm

Modern neurotechnology poses an unprecedented challenge to the *forum internum* because it can access mental processes without the individual's conscious control or awareness. The UN report identifies this as one of the defining characteristics of neurotechnology: it can "bypass the individual's conscious control or awareness" and "enable non-consensual external access to thoughts, emotions and mental states."

Even if current neurotechnology cannot read "actual" thoughts in the way depicted in science fiction, the UN report warns that these devices can "serve as 'mindreading-like' devices to make inferences about individuals." With the support of artificial intelligence, neurotechnologies can "elaborate sophisticated individualized psychological profiles and use them to predict individuals' behaviour or intentions or to infer their mental characteristics or predispositions."

This capacity for inference-based mindreading is already a reality. The UN report notes that "neuroimaging is increasingly explored and can be employed to make inferences about individuals on a large spectrum of mental states, including memories, semantic knowledge, emotions, dreams, inner speech and intentions." The *forum internum* is no longer a fortress that cannot be breached; it is becoming a landscape that can be mapped, analyzed, and exploited.

The Absolute Prohibition on Interference

The UN report is unequivocal about the implications of this technological capability: "Individuals have a right not to reveal their thoughts." This right is not subject to balancing against other interests. It is absolute. The report further warns that "some proposed applications of neurotechnologies in the area of justice and national security may lead to the criminalization of mere thought." People could be "punished for unexpressed thoughts or intentions" and "pre-emptive decisions affecting their lives or interest can be taken even before they translate into action."

This represents a fundamental departure from centuries of legal tradition, which has always required an act—not merely a thought—to trigger legal consequences. The *forum internum* has historically been protected precisely because it is inaccessible to external authority. Neurotechnology collapses this distinction, making the inner realm visible and actionable.

The Philosophical Foundation: Why Thoughts Are Sacred

The sacredness of the *forum internum* rests on a recognition that human dignity requires a space of absolute autonomy. The UN report grounds its analysis in this principle: "Neurotechnologies and their applications should first be examined against the foundational principle of human dignity, which encapsulates the essence of what it means to be human." This requires assessing "the capacity of neurotechnologies to affect the constitutive elements of the person: the individual's personhood, mental capacities and personality."

When neurotechnology allows external actors to read or alter mental processes, it treats the individual as an object rather than a subject. The UN report warns that this can lead to "the objectification and instrumentalization of individuals" and can "neglect autonomy by altering the foundation of their mental processes and states." The *forum internum* is sacred because it is the seat of agency—the place where individuals can think freely, form their own beliefs, and make autonomous choices.

The Practical Necessity of Protection

The UN report makes clear that protecting the *forum internum* is not merely a philosophical exercise but a practical necessity. The report notes that "neurotechnologies can be used to interfere and manipulate individuals" through "neuromodulation devices" that can alter "the physical and mental processes of a person's inner sphere in ways similar to 'brainwashing.'" Commercial applications like neuromarketing and neurogaming can "condition the forming of opinions, as well as influencing an individual's decision-making processes."

The report specifically warns that "the risk that such interference occurs even without the individual's consent or knowledge is high." Consumer neurotechnology devices are already being marketed for use during sleep, and immersive computing systems can "adjust experience on the basis of the neurodata detected" in ways that may be "misused for behavioural modification."

The sovereignty of thought is therefore not an abstract ideal—it is the foundation of democratic society, personal autonomy, and human dignity. Protecting the *forum internum* requires recognizing that the mind is the last private place, and that no technology, no commercial interest, and no state security apparatus should be permitted to breach its walls without the most rigorous legal and ethical safeguards. As the UN report concludes, "a coherent protection of the mind's inner space requires clarification of the interrelation between freedom of thought and other essentially interconnected rights," and the development of "specific standards" to make this ancient right relevant in the age of neurotechnology.

CHAPTER 20

The Philosophy of Mental Privacy — Locke to Bostrom

The legal and political case for neuro-rights rests on a deeper philosophical foundation. To understand why mental privacy is sacred, why cognitive liberty is the foundational freedom, and why the prohibition on non-consensual modulation of mental processes is absolute, one must engage with the long tradition of philosophical reflection on the nature of mind, person, and self. This chapter traces that tradition from the early modern philosophers who articulated the modern concept of personhood through to contemporary thinkers who confront the implications of neurotechnology head-on.

The Cartesian Inheritance

Modern philosophy of mind begins with René Descartes. In the Meditations on First Philosophy (1641), Descartes identified the inner life of thought as the one thing whose existence could not be doubted: cogito, ergo sum. The body, the world, the senses — all could be illusions. But thinking itself, the immediate experience of consciousness, could not be denied without contradiction. This insight grounded what came to be called the Cartesian dualism: a sharp distinction between the mental and the physical, between the realm of thought and the realm of extended matter.

The Cartesian view had profound legal implications. If the mind is in principle inaccessible to physical investigation, then the inner life of a person enjoys a natural privacy that no outside power can violate. Confessions extracted by torture were unreliable not only because they were morally objectionable but because the genuine contents of mind could never be confirmed by external means. The privilege against self-incrimination, the prohibition on coerced testimony, the doctrine that thought itself cannot be criminalized — all of these legal principles reflect, in one way or another, the Cartesian assumption that the inner forum is inaccessible.

Neurotechnology challenges this assumption directly. It does not refute Cartesian dualism as a philosophical doctrine; that argument continues. But it removes the practical basis on which the doctrine's legal implications rested. If thoughts can be inferred from brain activity, however imperfectly, then the inner forum is no longer inaccessible. The legal principles that depended on that inaccessibility now need a different grounding — one based not on the impossibility of access but on the impermissibility of it.

Locke and the Self as Continuity of Consciousness

John Locke, writing half a century after Descartes, offered a different but related account of the inner self. In the Essay Concerning Human Understanding (1690), Locke argued that personal identity consists not in the substance of the soul but in the continuity of consciousness. A person is the same person across time because their memories, perceptions, and reflections form a continuous chain of conscious experience. This account

had powerful legal and moral consequences. Responsibility attaches to the conscious self, which is the unit that experiences, remembers, and acts. Punishment is justified because it is administered to the same continuous self that committed the act.

Locke's account makes the question of mental sovereignty even more pressing in the neurotechnology age. If personal identity is constituted by the continuity of conscious experience, then technologies that alter that experience — that introduce false memories, that modulate emotional responses, that interrupt the chain of consciousness — strike at the foundation of identity itself. The harm of such technologies is not merely that they violate privacy or impose physical effects; the harm is that they fragment, distort, or impose discontinuities in the very self that the legal system holds responsible for action. A Lockean argument for neuro-rights is, at root, an argument that personal identity is a thing worth protecting against external interference, because without it, none of the other legal protections meaningfully apply.

Kant and Autonomy

The most enduring philosophical foundation for modern human rights is Immanuel Kant's principle of autonomy: the capacity of a rational agent to legislate moral law for themselves. For Kant, the dignity of the human person consists in this capacity for autonomous self-determination. Persons are not to be treated merely as means; they must be treated as ends in themselves, capable of acting on principles of their own rational endorsement.

Kant's framework grounds the central neuro-rights insight that the capacity for autonomous decision-making must be protected against external interference. A technology that influences a person's reasoning without their knowledge or consent treats that person merely as a means; it manipulates rather than persuades. The distinction between persuasion and manipulation, central to political and commercial ethics, becomes operational rather than rhetorical in the neurotechnology age. Persuasion engages the rational capacities of the audience; it operates through reasons. Manipulation bypasses those capacities; it operates through mechanisms the audience does not recognize and could not endorse if they did. Neurotechnology, particularly when applied without disclosure to its subject, falls naturally on the manipulation side of this line.

The Kantian argument provides the philosophical case for treating informed consent as the central operative concept in neurotechnology ethics, and for treating consent itself as something that can be undermined by certain kinds of intervention. A consent extracted by mechanisms that bypass rational deliberation is not consent in the morally significant sense, even if its formal indicia (a signature, a checkbox, a verbal yes) are present. This is why neuro-rights frameworks must distinguish between formal consent and substantive consent, and must build protections against the technological erosion of the latter.

Mill, Liberty, and the Sovereignty of the Individual

John Stuart Mill's On Liberty (1859) articulated, in language that has shaped liberal thought ever since, the principle of individual sovereignty: over himself, over his own body and mind, the individual is sovereign. Mill's argument was directed against the social pressures that

constrain individual development — the tyranny of custom, the despotism of public opinion. He was writing in a society that did not yet imagine the technological means by which individual sovereignty could be threatened. But his principle generalizes.

Cognitive liberty, as it has emerged in twenty-first-century scholarship, is the direct extension of Mill's principle to the cognitive domain. If individuals are sovereign over their bodies and minds, then they have the right to think their own thoughts, to use cognitive tools of their choosing, and to refuse cognitive interventions imposed by others. The Millian argument resists both compulsion and prohibition: a sovereign individual is one who can choose to use a cognitive-enhancing technology and one who can refuse to use a cognitive-monitoring technology. The state's role is to secure the conditions of choice, not to direct the choice itself.

This Millian frame has tensions that neuro-rights drafting must navigate. On one hand, it argues against paternalistic restrictions on adult use of neurotechnologies for their own purposes. On the other, it argues for strong protections against involuntary use of those same technologies. Drafting language that captures both sides of this commitment — enabling chosen uses while preventing imposed ones — is one of the central challenges of contemporary neuro-rights work.

Phenomenology and the First-Person Perspective

Twentieth-century phenomenologists — Husserl, Merleau-Ponty, Heidegger — emphasized that the inner life of consciousness has a structure that cannot be reduced to its physical substrate. The experience of being a self, of being present in one's body, of relating to the world from a particular point of view, is irreducibly first-person. Third-person description, however precise, leaves something out: the felt quality of experience, what philosophers call qualia.

The phenomenological tradition contributes a distinctive argument to the neuro-rights debate. Neurotechnology, particularly in its decoding applications, operates entirely in the third person: it reads physical signals and produces inferred descriptions. Even when those inferences are accurate, they are not the experiences themselves. To equate the inference with the experience is a category mistake. A legal regime that treats neural data as merely another form of biometric data — like fingerprints or facial geometry — misses what makes neural data distinctive. Neural data is not a measure of a physical trait; it is the physical correlate of subjective experience. To collect, process, and trade in it is to traffic in the substrate of conscious life. The phenomenological insistence on the irreducibility of the first-person perspective is, ultimately, what makes neural data different in kind from any other category of personal data.

Contemporary Philosophers and the Neuro-rights Frame

Contemporary philosophy has begun to engage with the neuro-rights agenda directly. Marcello Ienca and Roberto Andorno's 2017 paper proposing four neuro-rights — cognitive liberty, mental privacy, mental integrity, and psychological continuity — has been one of the most cited contributions, providing the conceptual scaffolding on which much subsequent

legal work has built. Their framework draws on Kantian, Millian, and Lockean elements to argue that the cognitive domain requires its own distinct rights vocabulary, not merely the application of existing rights.

Other contemporary thinkers have made related contributions. Nita Farahany's work on cognitive liberty connects the philosophical case for neuro-rights to the practical realities of consumer neurotechnology, arguing that the legal recognition of cognitive liberty is the most defensible basis for regulating technologies that affect the mind. Patricia Churchland and Daniel Dennett, working from very different positions in the philosophy of mind, have nonetheless contributed to public understanding of what neuroscience can and cannot reveal about the self. Nick Bostrom and the broader transhumanist tradition have argued for the importance of preserving cognitive autonomy in the face of enhancements they otherwise support — an interesting position that reveals neurotechnology to be as much a concern for those who embrace it as for those who fear it.

Why Philosophy Matters for Drafting

It would be possible to draft neuro-rights legislation without engaging the philosophical literature. Many drafters have done so. The resulting instruments tend to be technically competent but philosophically thin, capable of regulating obvious abuses but vulnerable to clever circumvention. The instruments that survive contact with sophisticated adversaries — corporate, governmental, or military — are those that articulate not merely what is prohibited but why. The why provides the interpretive resources that judges, regulators, and citizens use to extend the protection to circumstances the drafters did not foresee.

The philosophical tradition surveyed here provides those interpretive resources. Cartesian inaccessibility, Lockean continuity of consciousness, Kantian autonomy, Millian sovereignty, phenomenological first-person experience, contemporary cognitive liberty — each provides a different angle from which the central neuro-rights commitments can be defended. Drafters who understand the philosophical foundations can craft statutory text that bends in productive ways when applied to novel situations. Drafters who do not, cannot. The chapters that follow assume this philosophical literacy as a precondition for the substantive work they describe.

CHAPTER 21

Mental Privacy in the Age of Brain-Reading

The *forum internum*—the inner realm of thought—has been protected for centuries not because it was inviolable, but because it was *inaccessible*. No army could breach it. No government could surveil it. No corporation could mine it for profit. That era is ending. Neurotechnology, once confined to medical clinics and research laboratories, is now creating a direct gateway between the human brain and external devices. The question is no longer whether mental privacy is valuable, but whether it can survive.

The Gateway to the Mind

The United Nations report on neurotechnology and human rights (A/HRC/57/61) identifies the core threat with stark clarity: "The particularity of neurotechnologies is that they allow a direct connection to be established between the human brain and external devices and thus provide a gateway to interfere with mental and cognitive functions." This is not hyperbole. Unlike previous technologies that could observe behavior or infer mental states from external cues, neurotechnologies operate at the level of neural signals themselves.

The report distinguishes between two fundamental operations: "reading" and "writing." Neuroimaging devices can "detect neural signals and allow the interpretation of mental states and behaviors," while neuromodulation techniques can "alter and modulate the functioning of the brain" through inhibition or stimulation of targeted regions. Both capabilities threaten mental privacy, but in different ways. Reading exposes the contents of the mind; writing alters them.

Can Companies Read Your Thoughts?

The short answer is: not yet, but they are getting closer—and the gap between current capabilities and science fiction is narrowing faster than most people realize.

The UN report is careful to note that "most neurotechnologies are still being tested and do not provide accurate results in the decoding of people's thoughts, which means that 'mind-reading' or 'mind-steering' are not still achievable." However, the report immediately adds a critical caveat: "Neuroimaging is, however, increasingly explored and can be employed to make inferences about individuals on a large spectrum of mental states, including memories, semantic knowledge, emotions, dreams, inner speech and intentions."

This distinction is crucial. Current technology cannot read your thoughts the way a person reads a book—it cannot pluck a specific memory or a private fantasy from your brain and display it on a screen. But it *can* make statistical inferences about your mental states. With the support of artificial intelligence and machine learning, neurotechnology systems can "elaborate sophisticated individualized psychological profiles and use them to predict individuals' behavior or intentions or to infer their mental characteristics or predispositions."

The report warns that these capabilities mean neurotechnologies "can serve as 'mindreading-like' devices." They may not read your thoughts, but they can read *about* you—your emotional state, your level of attention, your subconscious reactions to stimuli, your likely future decisions. For commercial and surveillance purposes, that may be enough.

The Consumer Neurotechnology Boom

The threat to mental privacy is not coming from military laboratories or government intelligence agencies alone. It is coming from the consumer market. The UN report documents a rapidly growing industry that has seen investments surge "from $331 million to $7.3 billion in just 10 years." Nearly half of all worldwide patent applications in neurotechnology are registered in the United States.

The report describes how "non-invasive neurotechnologies" are being commercialized for "well-being, in the workplace or as an educational tool." Wearable devices track concentration levels, monitor sleep, and claim to improve cognitive performance. The new generation of wireless earbuds includes "biosensing electrodes, enabling the measurement of various biosignals, including brain activity." Users can check their brain activity on their phones as casually as they check their email.

But here is the critical finding: "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge. Moreover, research on potential side-effects, long-term impacts and their possible irreversible character is lacking." Consumers are handing over access to their neural data for products that may not even work as advertised.

The Data That Cannot Be Erased

The UN report identifies a fundamental asymmetry in the neurodata economy. Consumer neurotechnology companies are "already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future."

This is not paranoia. The report notes that neurodata "provide highly sensitive information and allow aspects of individuals' identity and personality to be revealed, such as sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions." Even if current decoding capabilities are limited, the data being collected today may be decryptable tomorrow. A person who agrees to have their brain activity recorded in 2025 may find that recording contains information they never intended to share when analysis tools improve in 2030.

The report warns that "with the wide deployment of brain-computer interfaces, requiring the collection, processing and storage of massive quantities of personal information, the risk of misuse of data, repurposing, predatory practices and hacking will become significant." Consent mechanisms, where they exist at all, are inadequate. The report finds that "data can be sold to third parties without the knowledge or authorization of users," and that users

"may be encouraged to exchange their data for the use of services on an (allegedly) voluntary basis."

The Bypassing of Conscious Awareness

Perhaps the most troubling aspect of neurotechnology's threat to mental privacy is that it can operate without the individual's knowledge. The UN report notes that neurotechnologies "bypass the individual's conscious control or awareness." A person wearing a consumer EEG headset may have no idea that their brain activity is being recorded, analyzed, and transmitted to a third party. A person using a neurofeedback app may not realize that their neural data is being used to train algorithms that will later be sold to advertisers or insurers.

The report specifically warns that "neurotechnologies allow the level of conscious reasoning to be bypassed, leaving individuals without psychological protection from having their mind involuntarily read." This is not a failure of consent—it is a structural feature of the technology. The brain does not have a "log out" button. It does not have a privacy setting. It generates signals constantly, whether the person is awake or asleep, aware or unaware.

The Limits of Existing Privacy Frameworks

The UN report is blunt about the inadequacy of current legal protections: "The specificities of brain data may not be adequately protected by existing international and domestic privacy protection frameworks." Data protection laws designed for credit card numbers and social media posts do not easily map onto neural signals that can reveal a person's innermost mental states.

The report calls for "context-specific standards to protect against the non-consensual intrusion of third parties" and warns that "consent requirements to avoid the unauthorized collection and processing of brain data should also be revisited and reinforced." The traditional model of informed consent—a checkbox clicked before using an app—is manifestly insufficient when the data being collected includes the electrical activity of a person's brain.

What Mental Privacy Requires

Mental privacy in the age of brain-reading requires more than just data protection. It requires recognizing that the mind is not just another source of personal information. As the UN report frames it, the right to privacy in this context must be "broadly approached and not restricted to the informational dimension." It must encompass the right to have one's inner mental life free from external intrusion, whether that intrusion takes the form of surveillance, manipulation, or coercion.

The report emphasizes the "interlinkage between the right to mental privacy and freedom of thought" as essential in this field. The two rights are not separate—they are two sides of the same coin. Freedom of thought is meaningless if thoughts can be read without consent. Mental privacy is meaningless if it can be overridden by commercial or state interests.

The sovereignty of thought, protected for centuries by the simple fact that no one could access it, now requires active legal and technical protection. The mind is no longer a fortress by default. It must be made one by design.

CHAPTER 22

Intellectual Property Rights to Thoughts — Can Companies Own Your Brain Data?

The question of who owns your brain data is not merely academic. As the UN Human Rights Council Advisory Committee noted in its 2024 report, neurotechnology companies are already "collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future." This is not a hypothetical concern. Consumer neurotechnology devices—headbands, earbuds, wristbands—are gathering neural signals today, and the companies that manufacture them are operating in what the UN describes as a "highly deregulated environment" characterized by "rapid expansion."

The core legal question is deceptively simple: Can a company own the electrical activity of your brain? The answer, under current law in most jurisdictions, is disturbingly unclear.

The Neurodata as Property Model

The prevailing approach in most countries treats neural data as a form of personal information—a subset of biometric data or health data subject to existing privacy frameworks. Under this model, companies that collect neurodata are data controllers, users are data subjects, and the relationship is governed by consent, purpose limitation, and data protection principles.

But this framework has a fundamental flaw. As the UN report emphasizes, "the specificities of brain data may not be adequately protected by existing international and domestic privacy protection frameworks." Neural data is not like other personal data. It is not something you choose to reveal, like a name or address. It is not something you generate through behavior, like browsing history. It is the direct electrical signature of your brain's activity—the physical substrate of thought itself.

Treating neurodata as property that can be owned, licensed, and sold raises profound questions about what it means to own a person's mental processes. The UN report warns that "neurodata provide highly sensitive information and allow aspects of individuals' identity and personality to be revealed, such as sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions." If a company owns this data, what rights does the individual retain over their own thoughts?

The Neurodata as Organ Model

An alternative framework, gaining traction among neuro-rights advocates, treats neural data more like an organ than a piece of property. You cannot sell your kidney on the open market in most countries. You cannot license your heart to a third party. The body is not property; it is the physical manifestation of the person. By extension, the argument goes, the electrical

activity of the brain—the functional expression of that organ—should not be treated as a commodity.

This approach has significant legal implications. If neural data is treated as part of the person rather than as property, then:

It cannot be owned by a company, only accessed with the individual's ongoing consent

It cannot be sold or transferred without explicit, informed, and revocable authorization

It cannot be used for purposes beyond those for which consent was given

It retains its connection to the individual even after collection

The UN report implicitly supports this view when it states that "the impact of neurotechnologies on privacy should be broadly approached and not restricted to the informational dimension." The report calls for "context-specific standards to protect against the non-consensual intrusion of third parties" and warns that "consent requirements to avoid the unauthorized collection and processing of brain data should also be revisited and reinforced."

The Commercial Reality

Despite these concerns, the commercial reality is that neurodata is already being collected, stored, and monetized. The UN report notes that "consumer neurotechnology companies are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future." The phrase "with the obvious expectation that they will become decipherable" is particularly significant. Companies are banking on the future ability to decode neural signals that are currently unreadable. They are stockpiling brain data today in the hope that tomorrow's algorithms will unlock its value.

The report also highlights the risk of "predatory practices" as "commercial interest in data collected through medical and non-medical neurotechnology devices will grow exponentially." With the neurotechnology market growing from $331 million to $7.3 billion in just ten years, the incentives to collect and monetize neural data are enormous.

The Patent Question

Beyond data ownership lies the question of whether neural patterns themselves can be patented. Can a company patent the specific neural signature associated with focus, relaxation, or emotional states? If so, does that give them ownership over the mental process itself?

The UN report does not directly address patent law, but its framework suggests caution. The report emphasizes that neurotechnologies "challenge the foundations of the human rights system" and that "respect for that overarching principle [of human dignity] should be monitored throughout the life cycle of neurotechnologies." Patenting neural patterns would effectively privatize aspects of human cognition, raising questions about whether such ownership is compatible with human dignity and the right to freedom of thought.

The Consent Paradox

Perhaps the most troubling aspect of the current legal landscape is what might be called the consent paradox. As the UN report notes, "it is hard to see how users can provide informed consent when their potential long-term effects are largely unknown." A user who agrees to share their neural data with a company today cannot know what that data will reveal in five or ten years, as decoding algorithms improve. They cannot know how the data might be used by third parties, or what inferences might be drawn from it.

The report emphasizes that "individuals may not fully understand the implications of using neurotechnologies" and "may unknowingly surrender control over their cognitive processes." This is not a failure of individual judgment; it is a structural feature of a technology that is advancing faster than our legal and ethical frameworks can adapt.

The UN report calls for "reinforcing standards on valid consent" and warns that "consent requirements to avoid the unauthorized collection and processing of brain data should also be revisited and reinforced." But the deeper question remains: Can consent ever be meaningful when the full implications of that consent are unknowable?

The Way Forward

The debate over whether neurodata is property or part of the person is not merely philosophical. It will determine the legal framework that governs the most intimate data ever collected from human beings. If neural data is treated as property, companies can own, license, and sell the electrical activity of your brain. If it is treated as part of the person, it receives the same protections as the body itself—and cannot be alienated or commodified.

The UN report does not prescribe a specific legal model, but its framework clearly points toward the organ model. The report's emphasis on human dignity, the absolute nature of freedom of thought, and the need to protect the *forum internum* suggests that neural data is not just another form of personal information. It is the physical expression of the inner self—and treating it as property risks reducing the person to a data source.

As the report states, neurotechnologies require assessing "whether their application contributes to the objectification and instrumentalization of individuals." The question of ownership is central to that assessment. If companies can own your brain data, they can own a piece of you. The law must decide whether that is acceptable—or whether some things are simply not for sale.

CHAPTER 23

Neurodiscrimination — When Your Brain Activity Determines Your Future

The UN report on neurotechnology and human rights (A/HRC/57/61) identifies a threat that is both subtle and profound: the use of neural data to discriminate. Unlike traditional forms of discrimination based on race, gender, or age, neurodiscrimination targets the very architecture of individual cognition—the unique patterns of thought, emotion, and predisposition that make each person who they are.

The report warns that "neuroprofiles may be used to discriminate among people on the basis of their mental signatures, leading to decisions (i.e. in relation to the workplace or insurance coverage) that may have a negative impact on the interests and rights of the data subject." This is not a distant possibility. The technology to create such profiles already exists, and the commercial incentives to use them are growing.

Employment: The Hiring Algorithm That Reads Your Mind

Consider the job interview of the near future. You sit across from a recruiter, but the real evaluation is being conducted by a non-invasive EEG headset that monitors your neural responses to questions. The device detects hesitation, emotional reactions, and subconscious associations. An algorithm compares your neural signature against a database of "ideal" employee profiles.

The UN report notes that consumer neurotechnology companies are already "collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future." What begins as a wellness device for concentration tracking becomes a tool for employment screening. You may never know that your brain activity was the reason you were rejected.

The report emphasizes that neurotechnologies "can serve as 'mindreading-like' devices to make inferences about individuals" and can "elaborate sophisticated individualized psychological profiles and use them to predict individuals' behaviour or intentions or to infer their mental characteristics or predispositions." In the employment context, this means that employers could reject candidates based on inferred traits such as susceptibility to stress, tendency toward creativity, or even political orientation—all without the candidate's knowledge or consent.

Insurance: Pricing the Risk of Your Thoughts

The insurance industry is built on risk assessment. Neurotechnology offers a new dimension of risk data that is more intimate than anything that came before. An insurance company could require policy applicants to undergo a brief neuroimaging session to assess their likelihood of developing neurological conditions, mental health disorders, or even risky behaviors.

The UN report specifically warns that "neuroprofiles may be used to discriminate among people on the basis of their mental signatures, leading to decisions (i.e. in relation to the workplace or insurance coverage) that may have a negative impact on the interests and rights of the data subject." This is not hypothetical. The same predictive algorithms that make neurotechnology powerful for medical diagnosis can be turned against individuals seeking coverage.

The report further notes that "neurodata provide highly sensitive information and allow aspects of individuals' identity and personality to be revealed, such as sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions." An insurance company could use this data to deny coverage or raise premiums based on neural markers that indicate a predisposition to depression, anxiety, or even antisocial behavior—regardless of whether the individual has ever exhibited such traits.

Education: Tracking the Developing Brain

Children are particularly vulnerable to neurodiscrimination. The UN report notes that "negative impacts on mental health of children, older persons and persons with disabilities should be considered together with the opportunities." In educational settings, neurotechnology could be used to monitor attention, track learning progress, and identify students who are "struggling" before they themselves are aware.

The report warns that neurotechnologies "may become pervasive throughout daily life despite the fact that, in most countries, applicable regulations are unclear, weak or non-existent." In the classroom, this means that a student's neural data could be used to stream them into different educational tracks, limit their access to advanced courses, or label them as having cognitive deficits—all based on inferences drawn from brain activity that may be inaccurate or biased.

The report specifically addresses the risk that "neurogaming devices can also drive compulsive or addictive use, particularly when neural reward circuits target altering consumption behaviours." For children, whose brains are still developing, the combination of neurotechnology and gamified learning could create feedback loops that shape their cognitive development in ways that are neither transparent nor reversible.

Criminal Justice: Punishing the Thought Before the Act

The most alarming form of neurodiscrimination may occur in the criminal justice system. The UN report warns that "some proposed applications of neurotechnologies in the area of justice and national security may lead to the criminalization of mere thought." With neurotechnology, "people can be punished for unexpressed thoughts or intentions and pre-emptive decisions affecting their lives or interest can be taken even before they translate into action."

This is not science fiction. The report notes that neurotechnologies "can be used to interfere and manipulate individuals" and that "through neuromodulation devices, the physical and mental processes of a person's inner sphere can be altered in ways similar to 'brainwashing.'"

In a criminal justice context, this could mean that individuals are subjected to "neurocorrectives"—coercive brain interventions aimed at altering their behavior or thought patterns.

The report is unequivocal: "No statements or confessions or other evidence obtained in violation of article 7 of the International Covenant on Civil and Political Rights may be invoked in any proceedings, not even during a state of emergency." Brain-reading techniques applied to detained persons are "clearly a form of prohibited pressure." Yet the report also notes that "the trend to introduce national security and law enforcement exceptions in artificial intelligence instruments and regulations raises concerns as they may not offer sufficient protection against problematic uses of neurotechnologies."

The Structural Bias Problem

Neurodiscrimination does not affect everyone equally. The UN report emphasizes that "groups experiencing structural discrimination and marginalized people or those in positions of socioeconomic disadvantage are the most at risk of misuse of neurotechnologies during the trial phase." Developers "may exploit their economic, physical or psychological vulnerabilities or those related to their work situation or age, among others."

This means that the same neural data that could be used to deny a job to a person with a mental health history could also be used to justify discriminatory policing practices. The report warns that "counter-terrorism policies" could use neurotechnologies "to pre-emptively punish non-violent action." In practice, this could mean that individuals from marginalized communities are disproportionately targeted for neural surveillance and pre-emptive intervention.

The report also notes that "persons with disabilities are particularly exposed to medical and scientific procedures, research and experimentation" and that "neurotechnologies are often developed without proper human rights guidelines and monitoring, and this may expose them to a disproportionate risk of physical and/or mental harms." Without meaningful participation of persons with disabilities in the design and development of neurotechnology, the risk is that these technologies "risk contributing to stigmatization and objectification, reinforcing structural-based discrimination."

The Legal Void

The UN report identifies a critical gap: "the specificities of brain data may not be adequately protected by existing international and domestic privacy protection frameworks." Most countries have laws against discrimination based on race, gender, religion, and disability. But no country has yet enacted comprehensive protections against neurodiscrimination.

The report calls for "context-specific standards to protect against the non-consensual intrusion of third parties" and warns that "existing loopholes in regulations, lack of technical expertise and capacity and the absence of adequate oversight bodies are factors that will certainly be exploited by large companies seeking profits."

The question is not whether neurodiscrimination will occur. The technology exists, the incentives are in place, and the legal framework is absent. The question is whether we will act in time to prevent the creation of a world where your brain activity—not your actions, not your choices, not your character—determines your future.

PART IV

Neurotechnology in Warfare and National Security

The battlefield of the 21st century is not defined by trenches, missiles, or even cyberattacks alone. It is defined by the human brain. As NATO's own doctrinal documents acknowledge, we have moved from "wars of attrition to wars of cognition." This is not a speculative future—it is the present reality, already being planned, funded, and in some cases, deployed.

The weaponization of neurotechnology represents a fundamental shift in the nature of conflict. Traditional warfare targets bodies and infrastructure. Cognitive warfare targets the mind itself—the processes of perception, decision-making, belief, and identity. As the NATO Science and Technology Organization's report on cognitive warfare states, it is "the art of using technologies to alter the cognition of human targets, most often without their knowledge and consent."

This is not psychological operations as traditionally understood. Cognitive warfare operates at a deeper level, exploiting the neural mechanisms that underpin thought itself. It feeds on disinformation and propaganda, yes, but it also draws on advances in neuroscience, neurotechnology, and artificial intelligence to achieve effects that were previously impossible. The goal is not merely to change what people think, but to alter how they think—to exploit cognitive biases, induce attentional tunneling, create cognitive stress, and ultimately disrupt the ability to make sound decisions.

The convergence of Nanotechnology, Biotechnology, Information Technology, and Cognitive Science (NBIC) has accelerated this transformation. What began as separate fields of inquiry are now converging into a unified capability to understand, predict, and manipulate human

cognition. The US Department of Defense formalized this convergence project in 2002, and major nations have since invested heavily in NBIC research for military applications.

China has been particularly explicit about its ambitions in this domain. The People's Liberation Army has defined "Military Brain Science" as a cutting-edge field that "can bring a series of fundamental changes to the concept of combat and combat methods, creating a whole new 'brain war' combat style and redefining the battlefield." China's cognitive warfare doctrine encompasses six categories of brain-related military functions: understanding the brain, protecting the brain, monitoring the brain, injuring the brain, interfering with the brain, and enhancing the brain. Each category corresponds to specific research and development programs, from sound and light weapons to brain-computer interfaces.

The implications are profound. As the NATO report notes, "Neuroscience is not subject to any rules of international law. A neuroweapon that attacks the brain is neither a biological nor a bioweapon." The Biological Weapons Convention and the Chemical Weapons Convention contain no provisions for neuroweapons. This legal vacuum creates a permissive environment for the development and potential deployment of technologies that target the most private space—the human mind.

The sections that follow examine the specific technologies, programs, and doctrines that constitute the emerging neuroweapons arsenal. From DARPA's Next-Generation Non-Surgical Neurotechnology program to the mysterious Havana Syndrome attacks, from directed-energy weapons to the concept of cognitive warfare itself, Part IV maps the terrain of a battlefield that exists not in any physical domain, but within the neural architecture of every human being.

CHAPTER 24

The History of Cognitive Warfare — From Propaganda to Neuroweapons

Cognitive warfare — the use of techniques and technologies to influence the cognition of populations and individuals for strategic purposes — is not new. The desire to influence what an adversary believes, perceives, and decides is as old as conflict itself. What is new is the technological capability to act on that desire with precision and at scale. The current generation of cognitive warfare grows out of a longer history of psychological operations, propaganda, and behaviour-modification research, and understanding that history is essential to recognizing the current frontier for what it is.

Ancient and Early Modern Origins

Sun Tzu, writing in the fourth century BCE, identified the supreme art of war as subduing the enemy without fighting — winning by altering the adversary's perception and will before any physical engagement. The practice of cognitive warfare in this broad sense has been continuous through the millennia. Romans, Byzantines, Mongols, Ottomans, and the European powers of the early modern period each developed techniques for influencing the morale of enemy soldiers, the loyalty of subject populations, and the perception of distant audiences. Religion, ritual, public spectacle, controlled information, and rumour were the principal instruments.

What changed with the modern era was scale. Printing, then mass literacy, then radio, then television, then the internet each enlarged the reach and the speed of cognitive operations. By the twentieth century, the techniques of propaganda had been systematized. Edward Bernays, the American public-relations pioneer and nephew of Sigmund Freud, articulated in his 1928 book Propaganda the proposition that the conscious and intelligent manipulation of the organized habits and opinions of the masses is an important element in democratic society. Bernays was describing the methods of commercial advertising, but the methods, once articulated, were rapidly adopted by states for political purposes.

The World Wars and the Industrialization of Influence

The First World War saw the first government-organized, industrial-scale propaganda campaigns. Britain's War Propaganda Bureau (Wellington House) and the United States' Committee on Public Information (the Creel Committee) developed techniques for mobilizing entire populations behind war aims, demonizing adversaries, and managing public response to wartime developments. The success of these campaigns left a lasting impression on observers, including on the rising authoritarian movements of the interwar period.

The Second World War saw the further refinement of these techniques. Nazi Germany under Joseph Goebbels developed propaganda into a comprehensive state instrument, integrated with film, radio, public ritual, and education. The Allied powers responded with their own

substantial programs, including the United States Office of War Information, the British Political Warfare Executive, and parallel agencies in the Soviet Union. The lessons learned in these efforts informed the development of postwar peacetime intelligence and counter-intelligence agencies and shaped the cognitive-warfare doctrines of the Cold War.

The Cold War and the Search for Mind Control

The Cold War saw the most ambitious and most ethically troubling phase of cognitive-warfare research to date. Both the Western powers and the Soviet bloc invested heavily in research programmes that sought to understand and ultimately to control human cognition. On the Western side, the most documented programmes were the CIA's MKULTRA (1953–1973) and its companion programmes ARTICHOKE and BLUEBIRD. These efforts examined chemical, biological, and psychological techniques for influencing behaviour, extracting information, and creating manipulable subjects.

MKULTRA's methods, as they have been declassified, included the administration of LSD and other psychoactive substances to unwitting subjects, prolonged sensory deprivation, hypnosis, electroconvulsive therapy, and combinations of these in various "deprogramming" or "reprogramming" regimes. The research was conducted at universities, hospitals, prisons, and clandestine sites in multiple countries. The most notorious episodes — the death of Frank Olson, the experiments at the Allan Memorial Institute under Donald Ewen Cameron, the involvement of unwitting bar patrons in Operation Midnight Climax — are well documented in the historical record. What is less documented is the full programme: many MKULTRA records were destroyed in 1973 on the orders of CIA Director Richard Helms, and what remains is fragmentary.

Soviet research, in parallel, examined chemical, electromagnetic, and psychological techniques for similar purposes. The term "psychotronics" entered the Soviet vocabulary to describe research into the effects of electromagnetic fields on consciousness and behaviour. The Soviet programme is less documented than the American one, in part because the relevant archives are less open and in part because Soviet defectors and researchers have published less extensively than their American counterparts. Enough is known, however, to confirm that the Soviet effort was extensive and that several of its lines of investigation continued under Russian successor agencies.

The Vietnam Era and the Refinement of Psychological Operations

The Vietnam War saw the United States develop psychological operations doctrine into a coherent discipline. The Joint U.S. Public Affairs Office (JUSPAO) and various military psychological-operations units developed leaflet, broadcast, and rumour-based techniques for influencing North Vietnamese and Viet Cong combatants and civilian populations. The results were mixed; the war was ultimately not won. But the doctrines refined during this period became the foundation of subsequent U.S. psychological-operations capacity, codified in successive Field Manuals and adopted by NATO allies.

The Soviet bloc developed parallel doctrines through its own institutions, with particular emphasis on active measures — deliberate operations to influence foreign perceptions

through forgery, fabricated press placements, recruitment of unwitting Western intermediaries, and amplification of authentic but selectively presented information. The KGB's Service A specialized in this work. The techniques developed during this period are recognizably the ancestors of contemporary disinformation operations conducted through social media.

Post-Cold War Continuities

The end of the Cold War did not end cognitive-warfare research; it merely reoriented it. The U.S. Department of Defense and intelligence community continued substantial investments in psychological-operations capability, information-operations doctrine, and behavioural-science research. DARPA, with its broader mandate, supported research on social-network analysis, persuasion modelling, and increasingly on direct neurological techniques. The Russian successor services rebranded but continued the active-measures tradition, applying it with renewed vigour as digital platforms emerged. China, building its own capabilities, developed doctrine that drew on both Soviet and Western traditions while integrating distinctively Chinese strategic concepts.

The watershed of the 2010s was the recognition that social media platforms had created an information environment in which sophisticated influence operations could be conducted at scale, at low cost, and with significant deniability. The Russian operations directed at the U.S. 2016 election, documented in the Mueller Report and subsequent investigations, were the public turning point. They demonstrated that state-directed cognitive operations conducted at scale could plausibly affect democratic outcomes. Similar operations have since been documented in dozens of countries.

The Neurotechnology Frontier

Contemporary cognitive warfare is converging with neurotechnology in several ways. First, social-media influence operations increasingly rely on inferred psychological profiles of their targets, derived from behavioural data and, in some cases, from physiological data including emotion-recognition systems. Second, research on direct neurological influence — transcranial stimulation, focused ultrasound, directed-energy techniques — has advanced from laboratory curiosities to operationally plausible capabilities. Third, the boundary between civilian and military neurotechnology research has eroded, with dual-use systems developed in academic and commercial contexts becoming available for security and military applications.

The combination of these trends — inferred psychological profiles, direct neurological techniques, and the erosion of civilian-military boundaries — represents a qualitative change in the nature of cognitive warfare. The earlier history was conducted in the medium of communication: messages were crafted, channels were employed, audiences were targeted. The new generation can, in principle, operate beneath the level of communication — directly on the neural substrate that processes messages, without the subject's awareness that anything has been communicated. This is not the cognitive warfare of leaflets and radio

broadcasts. It is something else, and the legal frameworks developed for the older form do not adequately cover the newer one.

The Continuity and the Discontinuity

Cognitive warfare exhibits both continuity and discontinuity. The continuity is in the goal: to influence cognition for strategic purposes. The discontinuity is in the technique: from communication to neural modulation, from message to substrate, from persuasion to manipulation in the strict philosophical sense. Recognizing both the continuity and the discontinuity is essential. Treating contemporary neurotechnology as a wholly new phenomenon, with no historical precedent, misses the institutional and doctrinal continuities that connect it to older traditions. Treating it as merely the latest iteration of an old form misses what is genuinely new about it.

The history matters for legal purposes because it identifies the institutions whose conduct most needs to be constrained. The agencies that conducted MKULTRA, the agencies that conducted active measures, the agencies that have been credibly implicated in directed-energy incidents — these are the institutions that have demonstrated, over decades, both the ambition to operate against the cognitive domain and the willingness to do so without democratic accountability. Constraining them is the institutional task to which neuro-rights frameworks must speak. Doing so requires the historical understanding that this chapter has attempted to provide.

CHAPTER 25

Cognitive Warfare — The Battle for the Brain

NATO defines cognitive warfare as "the art of using technologies to alter the cognition of human targets, most often without their knowledge and consent." This definition, proposed by Francois du Cluzel of NATO's Allied Command Transformation, captures the essence of a new form of conflict that targets not territory, infrastructure, or even information systems, but the human mind itself.

Cognitive warfare represents a fundamental shift in the nature of conflict. It is not merely an extension of psychological operations (PsyOps) or information warfare, though it draws on both. As the NATO Science and Technology Organization report makes clear, cognitive warfare "does not focus strictly on the field of 'information' but on that of 'cognition'—i.e., what the brain does with information." The distinction is crucial: information warfare targets the data we receive; cognitive warfare targets the neural processes by which we interpret that data.

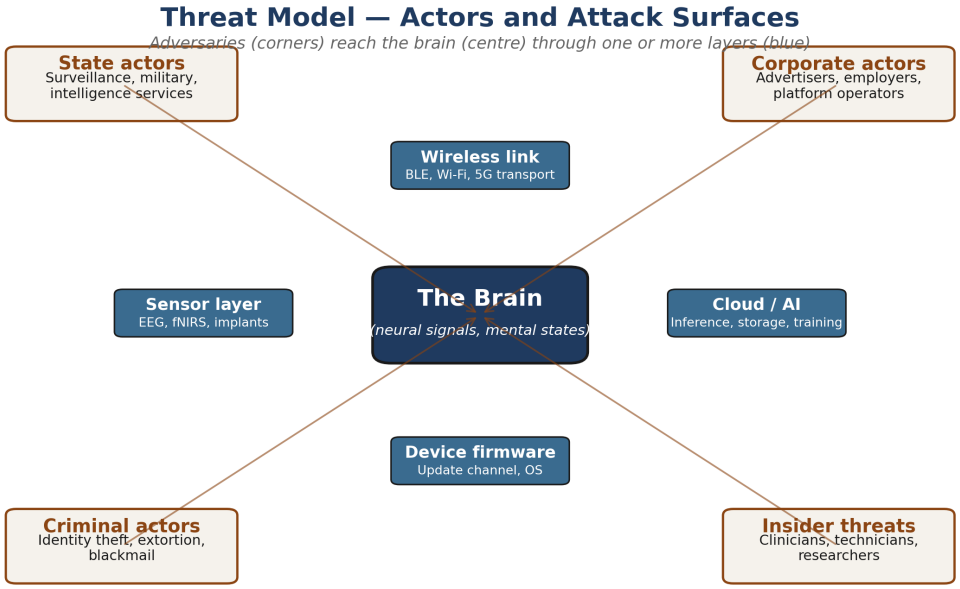


Figure 18.1. *Threat model for the inner realm. Adversaries (corners) reach the brain (centre) through one or more layers (blue). Effective defence must address every layer; no single point of intervention is sufficient.*

The Third Dimension of Conflict

Historically, warfare operated in two dimensions: the physical and the informational. Cognitive warfare introduces a third. To the physical domain of kinetic operations and the informational domain of data and communications is now added the cognitive domain—the realm of perception, decision-making, and belief formation.

The consequences of this expansion are profound. Unlike physical attacks, which are visible and attributable, cognitive attacks are "insidious" by nature. They "disrupt the ordinary

understandings and reactions to events in a gradual and subtle way, but with significant harmful effects over time." The target may never know they are under attack. Their decisions may be altered, their biases exploited, their trust eroded—all without a single shot fired or a single piece of malware detected.

Cognitive warfare has "universal reach, from the individual to states and multinational organisations." It feeds on the techniques of disinformation and propaganda, but goes far beyond them. Where traditional propaganda sought to change what people think, cognitive warfare seeks to change *how* people think—to alter the cognitive processes themselves.

The OODA Loop Under Attack

The military concept of the OODA loop—Observe, Orient, Decide, Act—provides a useful framework for understanding cognitive warfare's targets. The "Orient" phase, which involves making sense of observed information, is particularly vulnerable. As the NATO report notes, "Getting orientation wrong, no matter how well an actor can Observe, how quickly they can Decide, and how concisely they can Act, can nonetheless mean the actor is caught playing the wrong game."

Cognitive warfare aims to corrupt the orientation phase. By manipulating perceptions, inducing cognitive biases, and creating attentional saturation or "tunneling," an adversary can ensure that even the most capable military force makes catastrophic errors. The goal is not to destroy the enemy's capacity to fight, but to ensure they fight the wrong war entirely.

The Convergence of Technologies

What makes cognitive warfare possible now, when similar concepts have existed for decades, is the convergence of multiple technological trends. The "massive explosion of behavioural data made available by the advent of social media" has given researchers unprecedented insight into human cognition and decision-making. Advances in neuroscience, artificial intelligence, and machine learning have provided tools to exploit this understanding.

The NBIC convergence—Nanotechnology, Biotechnology, Information Technology, and Cognitive Science—has accelerated this transformation. What began as separate fields of inquiry are now converging into a unified capability to understand, predict, and manipulate human cognition. The US Department of Defense formalized this convergence project in 2002, and major nations have since invested heavily in NBIC research for military applications.

The Weaponization of Neuroscience

The NATO report identifies specific cognitive functions that can be affected by neurotechnology: memory, learning, and cognitive speed; wake-sleep cycles, fatigue, and alertness; impulse control; mood, anxiety, and self-perception; decision-making; trust and empathy; and movement and performance. In military settings, modifying these functions can be used to "mitigate aggression and foster cognitions and emotions of affiliation or

passivity; induce morbidity, disability or suffering; and 'neutralize' potential opponents or incur mortality."

These are not theoretical possibilities. The report states unequivocally that "the use of neuroscience and technology for military and intelligence purposes is realistic, and represents a clear and present concern." A 2014 US report confirmed that neuroscience and technology "had matured considerably and were being increasingly considered, and in some cases evaluated for operational use in security, intelligence, and defense operations."

The Chinese Doctrine of "Mind Superiority"

China has adopted an even broader definition of cognitive warfare that includes "the systematic utilisation of cognitive science and biotechnology to achieve the 'mind superiority.'" The People's Liberation Army has defined the "Cognitive Domain of Operations" as the battlefield for "conducting ideological penetration… aiming at destroying troop morale and cohesion."

China's Military Brain Science program is organized around nine functions: understanding the brain, protecting the brain, monitoring the brain, injuring the brain, interfering with the brain, repairing the brain, enhancing the brain, simulating the brain, and arming the brain. Each function corresponds to specific research programs and potential weapons systems.

Of particular concern is China's focus on "interfering with the brain"—causing "brain dysfunction and a loss of control with 'smokeless' methods." This language suggests development of directed-energy weapons or other technologies capable of affecting neural function without leaving physical traces.

The Legal Vacuum

Perhaps most troubling is the legal status of cognitive warfare and neuroweapons. As the NATO report notes, "Neuroscience is not subject to any rules of international law. A neuroweapon that attacks the brain is neither a biological nor a bioweapon." The Biological Weapons Convention and the Chemical Weapons Convention, the two existing treaties that might address such weapons, "do not contain any provisions for neuroweapons."

This legal vacuum creates a permissive environment for development and deployment. Cognitive attacks that alter decision-making, exploit biases, or disrupt neural function may not qualify as acts of war under international treaties. They exist in a gray zone below the threshold of armed conflict, where traditional rules of engagement do not apply.

A New Domain of Operations

The NATO report raises the question of whether the Alliance should declare cognitive warfare a new domain of operations, alongside land, sea, air, space, and cyberspace. The argument for doing so is that "a cognitive attack is not a threat that can be countered in the air, on land, at sea, in cyberspace, or in space. Rather, it may well be happening in any or all of these domains, for one simple reason: humans are the contested domain."

The human is both the vulnerability and the target. As the report concludes, “If kinetic power cannot defeat the enemy… psychology and related behavioural and social sciences stand to fill the void.” The battle for the brain has already begun, and the rules of engagement have yet to be written.

CHAPTER 26

Havana Syndrome — The First Salvo of Neurowar?

In late 2016, in Havana, Cuba, a group of U.S. diplomats, intelligence officials, and military personnel began experiencing mysterious and often debilitating symptoms reminiscent of traumatic brain injury—but without any precipitating trauma. Victims reported headaches, dizziness, fatigue, nausea, anxiety, vertigo, memory loss, and other cognitive difficulties. Most strikingly, all described feeling waves of pressure in their heads, ranging from dull discomfort to immediately overwhelming. Many noted that simply moving from one room to another alleviated the symptoms.

These incidents were not limited to government officials. Children and pets displayed peculiar and troubling behaviors. Yet many affected individuals did not immediately report their private health concerns, struggling to understand what was happening to them. Only in the past few years has this information surfaced in the public domain, and there have been no shortage of theories: sonic weapons, microwaves, insecticides, crickets, mass hysteria, and everyday aches and pains.

But the evidence points to something far more sinister.

A Pattern of Attacks

What became known in the media as “Havana Syndrome” is not an isolated phenomenon. Similar incidents have occurred in at least sixteen countries, including the United States. As CIA Director William Burns stated in July 2021, “There are probably a couple of hundred incidents across the U.S. government and across the globe.” The attacks have followed a disturbing geographic pattern:

Cuba (2016–2017): The initial cluster at the U.S. Embassy in Havana

Uzbekistan (2017): A similar incident at the U.S. Embassy in Tashkent

Russia (2017): Attacks reported in Moscow

China (2017–2018): Incidents at the U.S. Consulate in Guangzhou

Poland, Georgia, Australia, Taiwan (2019): Expanding geographic reach

United Kingdom and United States (2019): Attacks on home soil

United States (2020): Including incidents near the White House

Kyrgyzstan (2020), Austria and Germany (2021), Vietnam (2021), India (2021), Colombia (2021): Continued global pattern

The 2021 incident in Vietnam was particularly alarming: it delayed Vice President Kamala Harris’s flight, demonstrating that even the highest levels of security provide no protection.

The Medical Evidence

A Department of State–funded team of neuroscience experts investigated the Cuba cases. Dr. James Giordano, a neuropathologist and one of the appointed scientists, stated in his 2018 USSOCOM briefing: "This is intentional, this is directed, this seems to be a beta test of some type of a viable neuroweapon."

An extensive 2020 report from the National Academies of Sciences concluded that the most plausible cause is some form of directed energy. The report documented that victims experienced "the sudden onset of a perceived loud sound, sometimes described as screeching, chirping, clicking, or piercing, a sensation of intense pressure or vibration in the head, and pain in the ear or more diffusely in the head." Acute symptoms included dizziness, fatigue, impaired balance, headache, and impaired concentration—many of which became chronic years after initial onset.

The directional and location-specific details in the patients' histories are "unlike any disorder reported in the neurological or general medical literature." This unusual pattern provided the basis for increased suspicions of directed-energy weapons.

The U.S. Special Operations Command defined this cluster of symptoms as "unconventionally acquired brain injury" (UBI). The Secretary of Defense's September 2021 memorandum termed them "Anomalous Health Incidents" (AHI). Both definitions acknowledge that something is happening to these individuals, even if the precise mechanism remains officially unconfirmed.

The Neuroweapon Hypothesis

The attacks in Cuba appear to represent one of the first-known uses of a neuroweapon. As defined by scholar Armin Krishnan, neuroweapons are "weapons that specifically target the brain or the central nervous system in order to affect the targeted person's mental state, mental capacity and ultimately the person's behavior in a specific and predictable way."

The implications are profound. Not only are individual victims being systematically removed from their posts—with decades of institutional knowledge and operational experience lost—but the attacks are having strategic effects. When the attacks in Cuba occurred at a critical juncture in U.S.-Cuba relations, the State Department recalled a majority of personnel, significantly cooling previously warming ties between the two states.

The Strategic Context

The Naval Postgraduate School thesis *Neurowar Is Here!* (December 2021) by Joshua D. Gramm and Brian A. Branagan argues that Havana Syndrome represents the emergence of a new form of warfare. The authors note that "great power competition has returned to the forefront of international relations, as China and Russia seek to contest America's global leadership." In this context, "the battle for influence begins and ends in the human mind, where reality is perceived."

The thesis documents that both China and Russia are actively developing neurotechnology for military applications. China's research includes "interfering with the brain" through "smokeless" methods. Russia's work builds on decades of research into "reflexive control"

and active measures. Both nations clearly believe in the importance of this new technology and see the military and strategic potential it offers.

From a strategic standpoint, neuroweapons are ideally suited for gray zone operations—actions below the threshold of armed conflict that nonetheless achieve strategic effects. They leave no physical evidence visible to conventional forensics, are perpetrated remotely across jurisdictions, employ technologies that can be officially denied or classified, and create isolation through systematic gaslighting by institutions.

A Beta Test for Neurowar

The pattern of Havana Syndrome incidents suggests something more than isolated attacks. The geographic spread, the targeting of multiple individuals at different diplomatic posts, the specific and consistent symptoms—all point to a systematic testing of capabilities. As Dr. Giordano suggested, these appear to be beta tests of viable neuroweapons.

The implications for international law are stark. As the NATO report on cognitive warfare notes, "Neuroscience is not subject to any rules of international law. A neuroweapon that attacks the brain is neither a biological nor a bioweapon." The Biological Weapons Convention and the Chemical Weapons Convention contain no provisions for neuroweapons. This legal vacuum creates a permissive environment for development and deployment.

Victims of these attacks are left without recourse. They lack formal complaint procedures that recognize technological attacks. Investigative bodies lack technical expertise to verify neurotechnological harassment. Medical professionals misdiagnose technological torture as mental illness. Perpetrators operate with complete impunity behind corporate or state secrecy.

The question is no longer whether neuroweapons exist, but whether the international community will recognize them before they become normalized instruments of statecraft. Havana Syndrome may be the first salvo of neurowar—a warning shot that the battlefield of the mind has already been claimed, and the rules of engagement have yet to be written.

CHAPTER 27

DARPA's N3 Program — Non-Surgical Neurotechnology for the Military

If Havana Syndrome represents the shadowy edge of what neuroweapons might already be, DARPA's Next-Generation Non-Surgical Neurotechnology (N3) program represents the open, funded, and deliberate pursuit of neural interfaces for able-bodied soldiers. Where the former is a mystery, the latter is a matter of public record—complete with performance metrics, milestone schedules, and a call for proposals that explicitly invites universities, government labs, and private companies to build the technology that could close the loop between a soldier's brain and a weapon system.

The Problem DARPA Seeks to Solve

The Department of Defense has a straightforward problem: current neural interfaces require surgery. Implanted electrodes can provide high-fidelity read and write capabilities to the brain, but they are appropriate only for patients with medical conditions like paralysis or Parkinson's disease. Able-bodied military personnel cannot use them. The risks of infection, the permanence of implantation, and the sheer impracticality of battlefield surgery make existing brain-computer interfaces irrelevant for the warfighter.

DARPA's N3 program, as described by program manager Dr. Al Emondi in a 2018 Proposer's Day presentation, aims to solve this problem by creating "reliable neural interfaces without the need for surgery or implanted electrodes." The goal is to enable soldiers to "facilitate multitasking at the speed of thought" and "interface with smart decision aids." The strategic vision is not rehabilitation—it is enhancement. The target user is not a patient but a soldier who needs to process information faster, control systems remotely, and maintain cognitive superiority over adversaries.

Two Technical Tracks

The N3 program is organized into two parallel technical areas, each pursuing fundamentally different approaches to the same end.

Technical Area 1: Noninvasive seeks to develop external devices that can read from and write to the brain using ultrasound, light, radio frequency, or magnetic fields. The signal quality target is "neural ensemble activity"—meaning the system must be able to decode patterns of firing from populations of neurons, not just gross electrical signals from the scalp. The performance metrics are extraordinary: spatial resolution under 1 cubic millimeter, temporal resolution under 10 milliseconds, closed-loop system latency under 50 milliseconds, and the ability to control at least 6 degrees of freedom while receiving at least 6 categories of sensory feedback simultaneously. The system must operate continuously for at least 2 hours, maintain accuracy of 95% or greater, and cause no more than 1 degree Celsius rise in the tissue being read from or written to.

Technical Area 2: Minutely Invasive pursues a different path: injectable, ingestible, or intranasal delivery of nanotransducers—particles smaller than 1,000 times the size of a cell body—that can interact with individual neurons. These nanotransducers would be delivered through the bloodstream, cross the blood-brain barrier, and target specific cell types. The performance metrics here are even more ambitious: single-neuron resolution, spatial resolution under 50 cubic micrometers, control of at least 10 degrees of freedom, and cell-type specificity that allows excitatory or inhibitory control of stimulation. The integrated device must be no larger than 125 cubic centimeters.

The Three-Phase Structure

Both technical areas follow a three-phase development structure, with clear milestones and deliverables.

Phase I (12 months) requires benchtop demonstration of independent read and write capabilities. Teams must fabricate subcomponents, characterize their resolution and latency, develop system-level designs including link margin power budgets and channelization approaches, and demonstrate read and write through "skull and tissue like" mediums.

Phase II (21 months) moves to in vivo proof of concept in animals and humans. Teams must integrate read and write subcomponents into a single device, develop algorithms to decode neural ensemble activity and encode sensory feedback, and evaluate physical properties including crosstalk, signal-to-noise ratio, and safety parameters.

Phase III (48 months) requires human demonstration of a closed-loop prototype system in a Department of Defense-relevant task. The final deliverable is a field-ready system that can decode motor and cognitive signals, encode sensory feedback, and enable a soldier to simultaneously perform multiple tasks.

The Notional System

The N3 program envisions a complete closed-loop system with three components. The integrated device, worn externally or deployed internally via nanotransducers, acquires neural signals and transduces them into a sensor input modality. It also receives stimulation parameters from an external processing unit and broadcasts stimulation waveforms back into the brain. The processing unit amplifies, filters, sorts, and multiplexes the signals, transmits them to the DoD application, receives incoming signals from the application, encodes them, and transmits them back to the integrated device. The DoD application itself—the weapons system, the drone, the decision aid—receives the decoded neural commands and provides feedback signals to be written back into the soldier's brain.

This is not a monitoring system. It is a bidirectional, closed-loop neural interface designed for real-time control of military hardware at the speed of thought.

The Ethical Dimension

Notably, the N3 program's Broad Agency Announcement explicitly requires proposers to address "Ethical, Legal, and Societal Implications (ELSI)" of their proposed technology. This

requirement, while commendable, raises the question of whether any institutional ethics review can adequately address the implications of a technology designed to merge the human brain with weapons systems.

As Jonathan Moreno observed in his 2004 analysis of DARPA's neuroscience programs, "Discussions about ethics require the oxygen of transparency, precisely the item in short supply in national security matters." Moreno, who served on two presidential advisory commissions, noted that DARPA's track record includes the Saturn rocket, ground radar, the Stealth Fighter, the Predator missile, and the Internet. The agency has a demonstrated capacity to transform fundamental research into world-changing technology. Its turn toward the brain represents a logical extension of its mission—but one that raises questions our society "has barely articulated, let alone carefully addressed."

The Implications for Neuro-Rights

The N3 program is not a secret project. It is a publicly funded, openly described research initiative with published metrics, scheduled milestones, and a call for proposals that any qualified institution can answer. Its existence demonstrates that the United States government is actively investing in technology that would, if successful, give soldiers the ability to control weapons with their thoughts and receive sensory feedback directly into their brains.

This is not science fiction. The program's timeline called for human demonstrations by 2022. Whether those demonstrations occurred, and what they revealed, remains a matter of public record that citizens and legislators have a right to know.

The N3 program also demonstrates that the line between therapeutic and enhancement applications of neurotechnology is not merely blurry—it is deliberately being erased. The same technology that could one day help a paralyzed person control a prosthetic limb could also enable a soldier to pilot a drone with her mind. The same nanotransducers that could deliver drugs to treat epilepsy could be used to enhance cognitive performance or suppress fear.

DARPA's N3 program forces a question that the international community has not yet answered: if a government can build a non-surgical neural interface for its soldiers, what protects the rest of us from the same technology being used without our consent?

CHAPTER 28

Russia's Psychotronic Programmes — A Historical Survey

The Soviet Union, and its Russian successor state, conducted some of the most extensive and least documented research into the electromagnetic and acoustic manipulation of human cognition in the twentieth century. Western awareness of this research has been fragmentary, drawn from defector testimony, declassified intelligence assessments, occasional academic publications, and the cultural shadow of the term psychotronics itself. A comprehensive history of the Soviet and Russian programmes is not possible from open sources. What is possible, and what this chapter attempts, is a survey of what is reasonably known, what is plausibly suspected, and what remains uncertain.

The Origins: Bekhterev and Soviet Reflexology

The Soviet tradition of brain research traces back to the late tsarist and early Soviet periods, particularly to the work of Vladimir Bekhterev (1857–1927) and the school of reflexology he founded. Bekhterev's institute, founded in St. Petersburg in 1907 and reorganized after the Revolution, conducted extensive research on the physiology of the nervous system, the conditioning of behaviour, and the possibility of remote influence on the mind. Bekhterev's own death in 1927, shortly after he was reported to have diagnosed Stalin as a paranoiac, has been the subject of speculation; the official explanation was natural causes.

What Bekhterev's institute established was a tradition of brain research that combined neurophysiology, behaviourism, and an openness to investigating phenomena that Western institutions of the same period tended to dismiss as occult. This combination shaped the trajectory of Soviet research for decades. Where Western brain research generally retreated from claims of remote influence after the methodological failures of early-twentieth-century parapsychology, Soviet research continued to investigate such claims, with substantial state funding, through the entirety of the Soviet period.

Pavlov and the Conditioning Tradition

Ivan Pavlov's work on conditioned reflexes is, of course, internationally famous. What is less generally appreciated is the extent to which Pavlovian conditioning shaped subsequent Soviet research on the manipulation of behaviour. Where American behaviourism developed in academic psychology departments under the influence of Watson and Skinner, Soviet behaviourist research developed in close coordination with state security agencies that recognized in it a tool for the management of populations and the extraction of information from individuals.

Pavlovian techniques were adapted for use in interrogation, particularly during the Stalinist period, with results that historians of the Soviet show trials have documented in detail. The reliability of confessions extracted under such conditions was, of course, low; the political utility was nonetheless considerable. The institutional infrastructure built during this period,

including specialized institutes for studying the physiology of behaviour, continued in modified form throughout the Soviet era.

The Microwave Era

Soviet research on the biological effects of microwave radiation began in earnest in the 1950s and continued through the dissolution of the Soviet Union. Where Western standards focused on the thermal effects of microwave exposure — the heating of tissue at high power densities — Soviet researchers documented and reported a much broader range of non-thermal effects, including effects on the central nervous system. The Soviet exposure standards for microwave radiation, in consequence, were orders of magnitude more conservative than the corresponding Western standards, a divergence that persists in modified form to the present.

The most publicly documented Soviet operational use of microwaves was the irradiation of the U.S. Embassy in Moscow, which began in the 1950s and continued for decades. The Moscow Signal, as it came to be called, directed pulsed microwave radiation at the embassy building from sources outside. The purpose has never been definitively established. Possibilities discussed in the declassified record include disruption of embassy electronic surveillance equipment, induction of biological effects on embassy personnel, and pure intelligence-gathering through novel means. The U.S. response, the Pandora Project, attempted to study the effects on personnel and resulted in inconclusive findings that have themselves been the subject of subsequent declassification disputes.

Psychotronics as a Field

By the 1970s, Soviet research on what was then called psychotronics had developed into a recognizable field with dedicated institutes, journals, and conferences. The field encompassed research on the biophysical bases of consciousness, the effects of electromagnetic fields on mental function, and the possibility of remote influence on cognition and behaviour. Western intelligence assessments from the period, some now declassified, documented Soviet investments in this research at scales suggesting state interest in operational applications.

The exact nature of those applications is, in the open record, unclear. Defector testimony from the late Soviet and early post-Soviet period describes specific programmes and specific capabilities, but the reliability of individual defector accounts is variable. Academic publications by Soviet and Russian researchers describe basic-science findings without explicit operational claims. The combination of substantial investment, secrecy, and selective declassification creates a picture in which the existence of programmes can be confirmed but the specific outcomes cannot.

The Post-Soviet Transition

The dissolution of the Soviet Union in 1991 disrupted, but did not end, this research tradition. Some institutes lost funding and contracted; others survived in modified form under Russian Federation auspices. Personnel migrated to commercial, academic, and

security positions. Some research that had been classified in the Soviet period entered the academic literature in the 1990s, providing one of the major windows through which Western researchers learned about the scope of the earlier programmes.

Under the Putin administration, particularly after 2000, Russian state interest in technologies for the manipulation of consciousness has been renewed and, in some respects, expanded. Publications by researchers associated with Russian security services have addressed topics including the analysis of mass psychology, the use of electromagnetic and acoustic techniques in security contexts, and the integration of such techniques with conventional information operations. The Russian Federation's investment in active-measures capabilities through agencies such as the GRU and its psychological-operations units has been documented in extensive detail by Western intelligence services and independent investigators.

The Ukraine War and Cognitive Operations

The Russian invasion of Ukraine, beginning with the 2014 annexation of Crimea and escalating dramatically in 2022, has provided an extended demonstration of contemporary Russian cognitive-warfare doctrine. Russian operations against Ukrainian and Western audiences have combined conventional propaganda with targeted disinformation, social-media manipulation, and selective use of cyber tools. The scale and sophistication of these operations exceed those of any previous post-Cold War period and indicate substantial continuing investment in the underlying capabilities.

Reports of more direct neurological techniques used in the Ukraine conflict are unconfirmed in the open record. Ukrainian sources have reported incidents that could be consistent with directed-energy effects, but conclusive forensic evidence has not been publicly released. The intense activity of the conflict, the unreliability of contested-zone reporting, and the strategic value to all parties of selective disclosure make this an area where definitive judgements are not currently possible. What is reasonably clear is that Russian cognitive-warfare doctrine treats the cognitive domain as a legitimate field of operations and that whatever capabilities exist within that domain are being employed.

What Open Sources Can and Cannot Tell Us

Western analysts seeking to understand Russian capabilities in this area face several methodological constraints. Russian classification practices are stricter than Soviet ones in some respects; the public-facing scientific literature is more selective; defector flows are reduced from the late Cold War period. The principal sources are open-source intelligence (commercial satellite imagery, social-media monitoring, academic publications), defector and émigré testimony, allied-government intelligence-sharing, and inferences from observable Russian operational behaviour.

What these sources together support is a picture of substantial continuing Russian investment in cognitive-warfare capabilities, including both conventional psychological operations and more advanced electromagnetic and acoustic techniques. What they do not support is detailed claims about specific weapon systems or deployments. Anyone making

such claims publicly, whether to support a policy argument or to dramatize a security concern, should be evaluated against this evidentiary background. Honest analysis acknowledges what is known and what is not.

The neuro-rights significance of the Russian programme is that it represents a major state actor whose disposition to use neurotechnology for purposes the Western legal framework would consider prohibited is unambiguous. Russian doctrine treats the cognitive domain as contested terrain. Russian capabilities, while not fully visible, are clearly substantial. Russian willingness to operate outside the constraints that Western states accept is repeatedly demonstrated in adjacent domains. A neuro-rights framework that does not account for this reality — that frames the question solely in terms of corporate consumer products in liberal democracies — will fail to protect against the threats that the contemporary international system most clearly contains.

CHAPTER 29

China's Military Brain Science Programme

If Russia represents the older tradition of cognitive-domain research, China represents the contemporary frontier. Chinese investment in neuroscience and neurotechnology has, in the past decade, grown to a scale that is now comparable to American investment in the same fields and is in some respects more strategically directed. The integration of civilian and military research, the formal recognition of military brain science as a state strategic priority, and the willingness to publish about offensive applications make the Chinese programme the most ambitious of any contemporary state actor. This chapter examines what is known about it.

Strategic Framing

The People's Liberation Army's framing of military brain science is articulated in publicly available Chinese military publications. The framework, attributed primarily to researchers at the Academy of Military Medical Sciences and affiliated institutions, identifies four broad categories of application: arming the brain, interfering with the brain, injuring the brain, and reading the brain. Each category encompasses both defensive and offensive applications, and each is described in published Chinese sources with a frankness that Western military publications generally avoid.

Arming the brain refers to enhancing the cognitive performance of one's own soldiers — through pharmacological agents, neurostimulation, biofeedback, and prospective brain-computer interfaces. The objectives include accelerated learning, sustained attention, reduced fatigue, and improved decision-making under stress. Interfering with the brain refers to disrupting an adversary's cognitive function — through electromagnetic effects, disinformation operations targeted on psychological profiles, and prospective directed-energy techniques. Injuring the brain refers to weapons designed to produce neurological damage. Reading the brain refers to surveillance, lie detection, and the extraction of information from individuals and populations.

The strategic significance of this fourfold framing is its comprehensiveness. Western military neurotechnology programmes, when they exist as identifiable programmes at all, tend to emphasize enhancement of one's own forces and defence against adversaries. The Chinese framing explicitly includes offensive applications across multiple modalities. The publication of the framing is itself a strategic communication: it signals to other actors that China considers neurotechnology a domain in which competition is open and in which Chinese investment is substantial.

The China Brain Project

The China Brain Project, officially launched in 2016 and incorporated into successive Five-Year Plans, is one of the world's largest neuroscience research initiatives. Its three primary research thrusts — the neural basis of cognition, brain disease diagnosis and intervention,

and brain-inspired computing — mirror in part the U.S. BRAIN Initiative and the EU Human Brain Project. The funding scale and the coordination across academic, industrial, and defence sectors are, however, distinctively Chinese.

The Project's civilian outputs include extensive publication on the basic neuroscience of memory, attention, perception, and motor control, frequently in major international journals. The Project's affiliated institutions have produced advances in macaque-monkey models of human disease, in optogenetic techniques, in non-human primate brain-computer interfaces, and in computational models of cortical processing. The quality of the work has been steadily improving and is now competitive with the best produced elsewhere. The Project's military and security dimensions are less visible in the academic literature but are documented through Chinese-language military publications and through Western intelligence assessments.

Brain-Computer Interface Research

Chinese researchers have produced a substantial volume of work on brain-computer interfaces in both invasive and non-invasive modalities. Tsinghua University, Tianjin University, Zhejiang University, and several PLA-affiliated institutions have all produced BCI research at the international frontier. Applications include motor neuroprosthetics for disabled patients, brain-state monitoring for educational and workplace settings, and military applications including pilot control augmentation and drone-swarm management.

One specific Chinese capability that has attracted Western attention is research on multi-person brain-coordination systems — BCIs designed to integrate the neural states of multiple individuals for coordinated decision-making or control of complex systems. Publications by researchers at the National University of Defense Technology describe experimental systems in which multiple operators jointly control a remote system, with each operator's neural signals contributing to the integrated command. The military applications of such systems are evident; the privacy and rights implications are extensive.

The Consumer Data Layer

The strategic asset most distinctive to the Chinese programme is the volume of neural data that Chinese deployments are generating. Independent investigation by journalists and researchers has documented the use of consumer-grade EEG headbands in Chinese classrooms and workplaces at scale. Products marketed under brands including BrainCo and various Chinese domestic competitors have been deployed in primary schools, secondary schools, and corporate offices, generating millions of hours of recorded EEG paired with contextual labels (subject matter, time of day, perceived performance, behavioural outcomes).

The data generated by these deployments is, in principle, available to Chinese state security agencies under existing Chinese law, which gives the state broad rights of access to private commercial data. The aggregation of this data — a major neural data corpus paired with rich contextual labels — represents the largest training dataset in existence for the development of population-scale neural decoders. Western training sets, generated under more restrictive

consent regimes and with less context, are orders of magnitude smaller. To the extent that decoder quality scales with training-set size and label richness, the Chinese position is structurally advantageous.

Public criticism within China of specific deployments has led to the withdrawal of some products from particular markets. The underlying data, once generated, is not recalled. The strategic value of that data is permanent. The case represents, perhaps more than any other example, how the absence of effective rights protection produces concrete capability advantages for the state that allows the absence.

The Xinjiang Question

The most serious allegations concerning Chinese use of neurotechnology relate to the Xinjiang Uyghur Autonomous Region. The available evidence is incomplete; Xinjiang is closed to independent monitors. What is known has been assembled from leaked documents, satellite imagery, testimony from former detainees and Uyghurs in diaspora, and academic publications by Chinese researchers describing the techniques used in security and detention contexts.

These sources document a security-and-surveillance infrastructure of unusual comprehensiveness, deploying conventional technologies (facial recognition, biometric collection, mobile-device monitoring) at population scale. Allegations of neurological screening, EEG-based lie detection, and behavioural assessment using neuroimaging or related techniques have been made in some sources but are less consistently corroborated than the conventional-surveillance documentation. Chinese researchers affiliated with security ministries have published academic work on the neural correlates of categories the Chinese state has applied broadly within Xinjiang, including frameworks of “extremism” and “terrorist intent.” The use of such frameworks in mass-screening contexts, if confirmed, would represent the largest population-scale deployment of neurotechnology in security operations anywhere in the world.

The Xinjiang case has implications that extend beyond the specific situation. It tests whether the international system can develop accountability mechanisms for population-scale neurotechnology deployments. It tests whether the technology’s most concerning applications can be addressed through existing human-rights mechanisms or whether new instruments are required. And it provides, for any future treaty negotiation, a template of what neuro-rights frameworks must, at minimum, be able to prohibit.

Strategic Implications for the Neuro-rights Movement

The Chinese programme has implications for the neuro-rights movement that practitioners must internalize. It complicates international negotiations: a binding global instrument requires Chinese participation, and Chinese incentives to accept restrictions are limited. It creates capability asymmetries: large training datasets generated without effective consent produce decoders that operate more reliably than those trained on smaller, more restricted datasets. It exports both products and norms: Chinese neurotechnology and the legal frameworks accommodating it spread to jurisdictions across the Belt and Road and beyond.

These implications do not argue for emulating the Chinese model. They argue for recognizing it, building the analytic capacity to understand it, and ensuring that the alternative liberal-democratic model is articulated, defended, and demonstrably worth choosing. The neuro-rights movement is, at this level, part of a larger contest of governance models. Its success will depend not only on the soundness of its legal arguments but on its ability to demonstrate that the protections it proposes are compatible with technological progress, scientific advance, and the legitimate security interests of the states adopting them. That demonstration is the strategic project of the next decade.

CHAPTER 30

NATO and the Cognitive Domain

The North Atlantic Treaty Organization has, in recent years, formally incorporated the cognitive domain into its strategic framework alongside the traditional physical domains of land, sea, air, and space, and the more recently recognized cyber domain. This recognition reflects an alliance-level acknowledgement that influence over the cognition of military forces and civilian populations has become a significant and contested terrain in modern conflict. The integration of cognitive operations into NATO doctrine, planning, and exercises is consequential for how the alliance prepares for future conflict and for how the neuro-rights movement engages with NATO member states.

The Sixth Domain

NATO's recognition of the cognitive domain as a distinct operational space has been articulated through multiple alliance publications and exercises. The basic proposition is that contemporary adversaries — state and non-state alike — conduct operations in the cognitive domain that have strategic effects equivalent to operations in the physical domains. Influencing decision-makers, manipulating populations, undermining trust in institutions, and shaping perception of events all produce military and political effects that traditional defence frameworks did not adequately capture. The alliance's formal recognition of the cognitive domain is, at its core, an attempt to capture these effects within the same doctrinal framework used to address other operational domains.

The implications of this recognition are doctrinal, organizational, and operational. Doctrinally, NATO planning now incorporates cognitive-domain considerations at the same level as kinetic-domain considerations. Organizationally, the alliance has established or designated bodies responsible for cognitive-domain activities, including the NATO Strategic Communications Centre of Excellence in Riga and the NATO Cooperative Cyber Defence Centre of Excellence in Tallinn. Operationally, exercises now incorporate cognitive-domain scenarios, and field commanders are trained to consider cognitive effects in their planning.

Defensive and Offensive Postures

NATO doctrine in the cognitive domain emphasizes defensive purposes — protecting alliance populations and forces from adversary cognitive operations. This emphasis reflects both the alliance's strategic orientation as a defensive pact and the political realities of operating across a membership that includes democracies with strong public commitments to information freedom. The principal tools developed in support of this defensive posture include strategic communications, counter-disinformation analysis, resilience-building in target populations, and improved cyber-defensive postures.

The line between defensive and offensive cognitive operations is, however, harder to draw in this domain than in others. The same tools that detect and counter adversary disinformation can be repurposed to conduct disinformation. The same analytical frameworks that identify

vulnerabilities in adversary information environments can be used to exploit those vulnerabilities. NATO doctrine does not articulate offensive cognitive-domain capabilities in publicly available documents, but the existence of such capabilities at the individual member-state level is widely understood, and the alliance's capacity to coordinate them in a contingency is implicit in its general planning frameworks.

The Neurotechnology Layer

Most NATO cognitive-domain activity to date has addressed what might be called the information layer of cognition — the level at which messages are crafted, transmitted, and received. The neurotechnology layer — the level at which neural substrates are directly engaged — has received less explicit doctrinal treatment. This relative inattention does not reflect a judgement that the layer is unimportant; rather, it reflects the early stage of the technology and the political sensitivity of articulating offensive applications publicly.

Individual NATO member states have, however, conducted substantial research on neurotechnology applications relevant to military operations. The United States, the United Kingdom, France, Germany, and Italy each have research programmes investigating cognitive enhancement of soldiers, neural monitoring of personnel, and brain-computer interfaces for system control. The U.S. DARPA programmes, including the Next-Generation Non-Surgical Neurotechnology (N3) programme, are the most publicly documented. Equivalent programmes exist, in smaller form, in other allied jurisdictions. The alliance-level integration of these national programmes is limited but increasing.

Havana Syndrome and the Alliance

The series of incidents collectively known as Havana Syndrome — anomalous health incidents affecting diplomatic and intelligence personnel in multiple countries since 2016 — has been a focus of NATO concern, both because of its impact on allied personnel and because of what it suggests about adversary capabilities in the cognitive domain. The U.S. government has produced multiple assessments of the incidents, with varying conclusions about cause. Other allied governments have produced their own assessments. The alliance as such has not produced a unified judgement.

What the incidents have done is concentrate attention on directed-energy effects as a potential capability and as a category of threat to be defended against. NATO member states have invested in improved diplomatic-security postures, screening protocols for personnel deployed in higher-risk environments, and medical surveillance for personnel reporting symptoms. The implications for the legal framework are addressed elsewhere in this guide; the implications for NATO operational planning are that the alliance now treats directed-energy effects as a category of possible threat for which defensive countermeasures are being developed.

Cognitive Warfare Exercises

NATO and member-state militaries have, in recent years, conducted exercises explicitly focused on cognitive-warfare scenarios. These exercises typically simulate adversary

influence operations directed at member-state populations during a crisis, with players in defensive roles responsible for detection, attribution, and counter-narrative response. The exercises have surfaced systematic weaknesses in alliance preparation: difficulties in coordinating responses across member states with different media environments, challenges in maintaining message coherence under time pressure, vulnerability to well-prepared adversary operations that exploit pre-existing public divisions.

More recent exercises have incorporated elements that approach the neurotechnology frontier: scenarios involving simulated directed-energy incidents, scenarios involving alleged uses of cognitive-influence technologies, scenarios involving the integration of conventional disinformation with claims about technological influence. The exercises remain within the conventional information-operations frame — they do not explicitly simulate the use of operational neurotechnologies against alliance populations — but the boundary is being approached. The doctrinal and operational responses that emerge from these exercises will shape NATO posture in the years ahead.

Implications for Neuro-rights

NATO's engagement with the cognitive domain has several implications for the neuro-rights movement. First, the alliance is an interlocutor whose engagement is consequential. Many member states' security postures are aligned through NATO processes. Norms adopted by the alliance influence national positions across the membership. Civil-society organizations and academic centres working on neuro-rights should engage with NATO bodies, particularly the Strategic Communications Centre of Excellence and the alliance's emerging cognitive-domain coordination structures.

Second, NATO's defensive orientation creates a natural fit between alliance interests and neuro-rights protections. Defensive cognitive-domain doctrine emphasizes the protection of alliance populations from adversary operations. Neuro-rights frameworks provide legal tools for that protection. The alignment is imperfect — NATO operational concerns extend to areas neuro-rights frameworks do not address, and neuro-rights protections extend to areas NATO operational concerns do not reach — but the overlap is substantial. Articulating that overlap in concrete proposals is a productive direction for engagement.

Third, the alliance's reluctance to articulate offensive capabilities publicly creates space for civil-society and academic interlocutors to propose, on a fresh slate, what the alliance's stance toward offensive cognitive-domain capabilities should be. The argument for an alliance-wide prohibition on certain categories of offensive neurotechnology — directed-energy weapons targeting persons, neuro-modulation devices used against civilian populations — has not yet been definitively rejected in NATO contexts and is in some ways more politically feasible at alliance level than at individual state level. The opportunity is short but real, and the neuro-rights movement should engage it directly.

CHAPTER 31

The Neuroweapons Arsenal — What Exists and What Is Possible

The line between defensive technology and weapon is often crossed not by the device itself, but by the intent with which it is deployed. A device that stimulates the vagus nerve to treat epilepsy becomes a weapon when it is used to induce seizures. A system designed to help a paralyzed person communicate becomes a weapon when it is used to implant commands. The neuroweapons arsenal, as it exists today, is not composed of exotic devices found only in classified facilities. It is composed of technologies that have been developed openly, for legitimate purposes, and then adapted, refined, or redirected toward hostile applications.

Directed-Energy Weapons: The Active Denial System

The most publicly acknowledged directed-energy weapon is the Active Denial System (ADS), developed by the U.S. Department of Defense. The ADS emits a focused beam of millimeter-wave electromagnetic radiation at 95 gigahertz. When it strikes a human target, the energy penetrates less than 1/64th of an inch of skin, heating the water molecules in the epidermis to a level that causes intense, intolerable pain. The pain ceases the moment the target moves out of the beam or the operator releases the trigger.

The ADS is classified as a “non-lethal” weapon. It is designed to disperse crowds, protect military installations, and provide an alternative to lethal force. But the technology itself is neutral: the same millimeter-wave emitter that causes pain could, with modifications to power, frequency, or modulation, be used to affect deeper tissues, including neural tissue. The question is not whether the physics permits such effects—it does—but whether the engineering has been pursued.

Microwave Hearing: The Foundation for Voice-to-Skull

The microwave hearing effect, also known as the Frey effect after its discoverer Allan H. Frey in 1961, is a well-documented phenomenon in which pulsed microwave radiation is perceived as sounds—clicks, buzzes, or even intelligible speech—within the skull. The effect occurs when microwave pulses cause rapid thermoelastic expansion in brain tissue, generating acoustic waves that are detected by the cochlea.

As V. N. Binhi notes in *Electromagnetic Mind Control: Fact or Fiction?*, “one can in principle modulate a high frequency signal with low frequency patterns that correspond to speech, and even electronically create certain statements.” This is not speculation; it is a direct extrapolation from established physics. The microwave hearing effect has been replicated in multiple laboratories. The U.S. Air Force studied it as a potential communication technology. The Soviet Union investigated it as a potential weapon.

The implications for neuroweapons are straightforward: if speech can be induced in a target’s head without their consent, then harassment, disorientation, command

implantation, and psychological torture become possible at a distance, without physical contact, and without leaving forensic evidence visible to conventional examination.

Electromagnetic Behavioral Modification: The Research Base

The scientific foundation for electromagnetic behavioral modification is not hypothetical. It has been demonstrated in controlled laboratory experiments, many of them funded by military agencies. The seminal calcium efflux experiments by W. Ross Adey and Susan Bawin in the 1970s showed that radiofrequency electromagnetic fields modulated at specific frequencies could alter calcium transport in brain tissue. This was not a thermal effect—the field intensities were too low to cause heating. It was a nonthermal interaction that opened the door to a new understanding of how electromagnetic fields can affect neural function.

Subsequent experiments by Dr. John Thomas and Dr. Abraham Liboff at the Naval Medical Research Institute in Bethesda found that rats exposed to specific low-frequency magnetic field combinations "temporarily lose their short-term memory after just one hour of exposure." Further research showed that "rats exposed to these magnetic field combinations can be made to learn either more quickly than controls or less, depending on the field combination that is used." The same fields could produce "greater aggressiveness or greater passivity."

These are not marginal effects. They are reproducible, systematic, and linked to specific biophysical mechanisms—particularly the NMDA glutamate receptor, which is central to memory and learning. As Liboff states, "the way has been pointed towards selective frequency control of behavior."

The Psychotronic Weapons Concept

The term "psychotronic weapons" originated in the Soviet Union and refers to technologies designed to influence human consciousness through electromagnetic or other physical means. As Binhi explains, "in the most extreme cases, according to conspiracy theorists, even simple innocuous laboratory EM generators became psychotronic generators that could be operated by a criminally motivated individual and change a person's state of mind or behavior."

But the line between conspiracy theory and classified reality is difficult to draw. Dr. Robert O. Becker, a pioneering researcher in electromagnetic medicine, "always maintained that both the US and Russian governments were very much involved in such research." He also believed there was "a deliberate attempt by authorities to cover up these clandestine activities by means of deliberate disinformation tactics"—including "weird and bizarre stories that were leaked to the public for one purpose: to muddy the waters, to make the whole question of mind control seem absolutely unbelievable."

The HAARP Question

The High Frequency Active Auroral Research Program (HAARP) in Alaska is a facility consisting of 180 antennas organized in a phased array, capable of focusing megawatts of

electromagnetic power into the ionosphere. Officially, HAARP is designed to study ionospheric processes and their effects on communication and surveillance systems. But the same phased-array technology that can beam energy into the sky could, in principle, be used to beam energy toward ground targets.

Binhi notes that "possible influences on the biosphere are widely debated." The physics is not in dispute: a phased array can focus electromagnetic energy on a specific location. Whether HAARP has ever been used for this purpose is a matter of public debate, not scientific impossibility.

The Weaponization Continuum

The neuroweapons arsenal is not a collection of discrete devices. It is a continuum that runs from known, declassified systems to speculative but physically plausible technologies. At one end are systems like the Active Denial System, whose effects are thermal and superficial. At the other end are systems that use nonthermal, frequency-specific electromagnetic fields to alter neural function—memory, learning, mood, aggression.

In between lie the microwave hearing effect, which can deliver auditory information directly into the brain; the potential for remote neural monitoring through detection of electromagnetic signatures; and the possibility of synthetic telepathy—two-way communication between machines and brains without the subject's awareness or consent.

As the UN Human Rights Council Advisory Committee has noted, "there are no discrete human rights policies to limit the misuse and abuse of neurotechnology." This regulatory vacuum exists precisely because the technologies themselves exist—in laboratories, in military programs, in patents, and, according to victim testimony, in operational deployment against civilian populations.

What Is Possible

The question is no longer whether neuroweapons are possible. The question is which of the possibilities have been realized, by whom, and against whom. The scientific literature establishes that electromagnetic fields can affect neural function, that specific frequencies can produce specific behavioral changes, and that these effects can be delivered remotely and covertly. The engineering literature establishes that the necessary power sources, antennas, and modulation systems exist. The historical record establishes that both the United States and the Soviet Union invested heavily in this research.

What remains unknown—because it remains classified—is the current operational status of these technologies. But the trajectory is clear. As Liboff writes in his foreword to Binhi's book: "I certainly think that governments have been actively involved for some time in trying to bring this possibility into reality. The underlying research is likely focused on two very different frequency ranges, first, at extremely low frequencies, from just one or two to several hundred hertz, and second, in the gigahertz region and beyond."

The neuroweapons arsenal is not a future threat. It is a present reality, hidden in plain sight behind the legitimate language of non-lethal weapons, crowd control, and defense research.

The question for the international community is whether to acknowledge this reality and build the legal frameworks to constrain it, or to continue pretending that the human mind remains beyond the reach of technological coercion.

CHAPTER 32

Hot War vs Cold War — How Neuroweapons Change Conflict

The distinction between hot war and cold war has always been blurry at the edges, but neuroweapons threaten to erase the boundary entirely. A weapon that targets the brain rather than the body does not respect the traditional thresholds that have defined armed conflict for centuries. The same technology that could disable an enemy battalion on a conventional battlefield could, with minimal modification, be used to influence a foreign election, silence a journalist, or slowly degrade the cognitive function of a diplomat stationed abroad.

Understanding how neuroweapons change conflict requires understanding how they would be used differently in these two contexts—and, more troublingly, how they enable operations that fall into neither category.

Hot War: The Battlefield

In a conventional hot war scenario, neuroweapons offer what military planners call "effects without casualties." The NATO concept of cognitive warfare, as articulated by Francois du Cluzel of Allied Command Transformation, defines this as "the art of using technologies to alter the cognition of human targets, most often without their knowledge and consent." On the battlefield, this could mean disabling enemy soldiers through directed energy that induces confusion, disorientation, or loss of consciousness—temporarily neutralizing combatants without killing them.

The Chinese People's Liberation Army has formalized this approach in its concept of "Military Brain Science," which explicitly includes categories for "injuring the brain" and "interfering with the brain." According to Chinese doctrine, these capabilities involve "promoting the research and development of sound, light, explosion, magnetic and other new types of weapons" and "causing brain dysfunction and a loss of control with 'smokeless' methods." The goal is not merely to kill but to disrupt decision-making at its source.

DARPA's N3 program aims to create non-surgical neural interfaces for able-bodied soldiers, enabling direct brain-to-machine communication on the battlefield. The logical extension of this capability is the ability to not only read but also write to the brain—to send commands, suppress fear responses, or induce specific emotional states. In a hot war, such technologies could give one side a decisive advantage: faster reaction times, better situational awareness, and soldiers who can be remotely controlled or calmed under fire.

But the same technologies that enhance friendly forces can be turned against them. As du Cluzel notes, "neuroscience is not subject to any rules of international law. A neuroweapon that attacks the brain is neither a biological nor a bioweapon." The Biological Weapons Convention and Chemical Weapons Convention, drafted before these technologies emerged,

contain no provisions for neuroweapons. This legal vacuum means that on a hot battlefield, there are no clear rules of engagement for brain-targeting weapons.

Cold War: The Covert Domain

The cold war use of neuroweapons is arguably more dangerous because it is more deniable. In this context, neuroweapons become tools of espionage, subversion, and gray zone conflict—operations that fall below the threshold of armed conflict but achieve strategic effects.

The Havana Syndrome attacks on U.S. diplomats from 2016 onward exemplify this model. Whether or not directed energy was the cause, the pattern of attacks—targeting individuals in their homes and hotel rooms, causing lasting neurological damage, and leaving no conventional forensic evidence—represents a template for how neuroweapons could be used in covert operations. The attacker achieves their objective (disrupting diplomatic operations, intimidating personnel, signaling capability) without triggering a military response, because the attack cannot be definitively attributed.

NATO's cognitive warfare doctrine explicitly acknowledges this gray zone application. As du Cluzel writes, "cognitive warfare has universal reach, from the individual to states and multinational organisations. It feeds on the techniques of disinformation and propaganda aimed at psychologically exhausting the receptors of information." When combined with actual neurotechnology—the ability to directly influence cognition rather than merely manipulate information—this becomes a weapon of unprecedented subtlety and power.

The Chinese concept of the "Cognitive Domain of Operations" treats this as a permanent state of conflict, spanning "from peacetime public opinion to wartime decision-making." The goal is "mind superiority"—not just defeating an enemy in battle but controlling how they think, what they believe, and how they make decisions. This is not war as traditionally understood; it is a continuous campaign of cognitive influence that never stops and never declares itself.

The Blurred Line

The most dangerous aspect of neuroweapons is that they make the line between hot war and cold war nearly impossible to draw. A directed energy device that causes neurological damage operates the same way whether it is used against a soldier on a battlefield or a civilian in their home. The difference is entirely contextual—and context is easy to dispute.

Victims of covert neurotechnology use, such as those who identify as Targeted Individuals, report patterns of harassment that include synthetic telepathy, remote neural monitoring, and electromagnetic manipulation. These reports, documented in submissions to the UN Human Rights Council, describe sustained campaigns of technological targeting that cause real suffering but leave no evidence that conventional forensic methods can detect. The perpetrators, whether state actors or non-state groups, operate with complete impunity behind the veil of deniability.

As du Cluzel notes, “non-kinetic actions are often elements of a global, discrete, even invisible action calling for an awareness of their danger and the use of effective defence and threat techniques to counter them.” The challenge is that the same technology can be a legitimate military tool in one context and a human rights violation in another—and there is no international framework to distinguish between the two.

The Strategic Implications

The weaponization of neuroscience changes the fundamental calculus of conflict. In hot war, it offers the possibility of disabling rather than killing, which could reduce casualties—but also lowers the threshold for using force, since the ethical and legal barriers to a “non-lethal” weapon are lower than for a lethal one. In cold war, it enables operations that are effectively invisible, leaving victims with no proof and no recourse.

The international community faces a choice: either develop frameworks that address neuroweapons across the full spectrum of conflict, or watch as the distinction between war and peace, combatant and civilian, legitimate military action and covert human rights abuse, becomes meaningless. The human mind, as du Cluzel writes, is “the contested domain”—and it is being contested right now, in ways that existing laws and institutions are not equipped to recognize or respond to.

CHAPTER 33

Where Is the Frontline? Defining the Battlefield of the Mind

If neurotechnology is used as a weapon, where do we draw the line? The question is deceptively simple. In traditional armed conflict, the frontline is a physical boundary—a line on a map, a trench, a river, a set of coordinates where opposing forces meet. But cognitive warfare has no such geography. Its battlefield is the human brain, and its frontline exists wherever a target thinks, decides, or believes.

The Challenge of Definition

NATO's own analysts acknowledge that there is "no commonly accepted definition of Cognitive Warfare." The concept remains poorly understood and is frequently confused with information warfare or cyber warfare. Francois du Cluzel, writing for NATO's Allied Command Transformation, offers a working definition: cognitive warfare is "the art of using technologies to alter the cognition of human targets, most often without their knowledge and consent."

This definition immediately raises the frontline question. If the target does not know they are being attacked, where does the attack begin? Where does it end? And how do we distinguish between legitimate influence operations—propaganda, psychological operations, public diplomacy—and illegal cognitive attacks that violate the sovereignty of the mind?

The source material from the Naval Postgraduate School thesis *Neurowar Is Here!* (Gramm & Branagan, 2021) frames the problem in stark terms: "The battle for influence begins and ends in the human mind, where reality is perceived." This is not a metaphor. When neurotechnology enables direct manipulation of neural processes—altering memory, emotion, decision-making, or perception—the battlefield is literally inside the skull of the target.

The Spectrum of Cognitive Attack

Du Cluzel identifies a spectrum of cognitive warfare operations, ranging from disinformation and propaganda to direct neurotechnological attack. The Chinese People's Liberation Army has formalized this spectrum in its concept of "Military Brain Science," which explicitly catalogs capabilities across ten categories: understanding the brain, protecting the brain, monitoring the brain, injuring the brain, interfering with the brain, repairing the brain, enhancing the brain, simulating the brain, and arming the brain.

The Chinese framework is instructive because it acknowledges no bright line between defensive and offensive applications. "Interfering with the brain" is listed alongside "protecting the brain." "Injuring the brain" sits next to "repairing the brain." The same research, the same technologies, the same institutions serve both purposes. This dual-use reality makes the frontline inherently ambiguous.

The Legal Vacuum

The most troubling aspect of this ambiguity is the legal vacuum it creates. As du Cluzel notes, "Neuroscience is not subject to any rules of international law. A neuroweapon that attacks the brain is neither a biological nor a bioweapon." The Biological Weapons Convention and the Chemical Weapons Convention were drafted before neuroweapons were conceivable, and they contain no provisions for them.

This means that a weapon that directly attacks the human brain—altering cognition, inducing fear, disrupting decision-making, or causing neurological damage—can be deployed without violating any existing international treaty. The attacker operates in a legal gray zone where the definition of "attack" itself is contested.

Gramm and Branagan argue that this is not a hypothetical concern. The mysterious "Havana Syndrome" attacks on U.S. diplomats, beginning in 2016, may represent "the first salvo of neurowar." If so, the frontline has already been crossed—but no one can agree on where it was drawn.

Proportionality and Attribution

Even if we could define the frontline, the principles of proportionality and attribution—central to just war theory and international humanitarian law—become nearly impossible to apply in cognitive warfare. How do you measure proportional response to an attack that leaves no physical damage, no casualties, no observable effects beyond altered cognition? How do you attribute an attack when the weapon is a directed energy beam or a targeted neural signal that leaves no forensic trace?

Du Cluzel observes that cognitive warfare "does not focus strictly on the field of 'information' but on that of 'cognition,' i.e. what the brain does with information." The cognitive effect is not a consequence of the action—it *is* the action. This makes traditional models of attack and defense inadequate. The OODA loop (Observe, Orient, Decide, Act) that governs military decision-making is itself the target. If an adversary can alter the "Orient" phase—changing how you perceive and interpret reality—then everything that follows is compromised.

The Human Domain

The frontline question ultimately forces us to recognize that the human being is not merely a participant in conflict but the contested domain itself. Du Cluzel quotes August Cole and Hervé Le Guyader: "the Human Domain is the one defining us as individuals and structuring our societies. It has its own specific complexity compared to other domains."

This recognition has profound implications. If the human mind is the battlefield, then every human being is a potential combatant—not because they carry a weapon, but because they possess a brain. Civilians, diplomats, journalists, political leaders, and ordinary citizens all become legitimate targets in cognitive warfare. The distinction between combatant and non-combatant, already strained in asymmetric warfare, collapses entirely.

The Chinese military brain science framework makes no distinction between military and civilian targets. Its goal is “mind superiority” achieved through “systematic utilisation of cognitive science and biotechnology.” The target is cognition itself, wherever it occurs.

Drawing the Line

Where, then, should the frontline be drawn? The answer must begin with a fundamental principle: the human mind is inviolable. No technology, no weapon, no operation should be permitted to alter cognition without the conscious, informed consent of the individual. This is the core of cognitive liberty—the right to control one’s own consciousness.

But this principle is easier to state than to enforce. The technologies that enable cognitive warfare are dual-use, often indistinguishable from legitimate medical or commercial applications. The attacks are covert, leaving no evidence that conventional forensic methods can detect. The perpetrators operate with deniability, exploiting the legal vacuum.

The frontline, in the end, is not a line on a map. It is the boundary between voluntary and involuntary influence on the mind. And that boundary, in the absence of clear legal frameworks and detection capabilities, is being crossed every day—without anyone knowing where the battle is being fought.

CHAPTER 33-bis

Where Mind-Control Risk Comes From — Vectors, Actors, Consequences

The book has so far treated the risks of unauthorised neural intrusion in dedicated chapters: cognitive warfare (Part IV), the targeted-individual phenomenon (Part V), workplace surveillance (Chapter 60), and so on. This chapter consolidates and extends. It is a map of where mind-control risk actually comes from in 2026 and over the next decade — which actors, which technologies, which routes — and what the consequences will be if the risk vectors are not closed.

The word "mind control" is loaded, and we use it with deliberate breadth. It does not mean only the cinematic scenario of remote, irresistible control of a person's behaviour by an external operator. It means any technological capability that materially impairs a person's ability to form, hold or act on their own thoughts, beliefs and decisions free of external coercion or manipulation that bypasses their conscious will. By that definition, mind-control risk is not a future possibility — it is a present condition, in some forms widely deployed.

Vectors of Mind-Control Risk — Where the Threats Converge

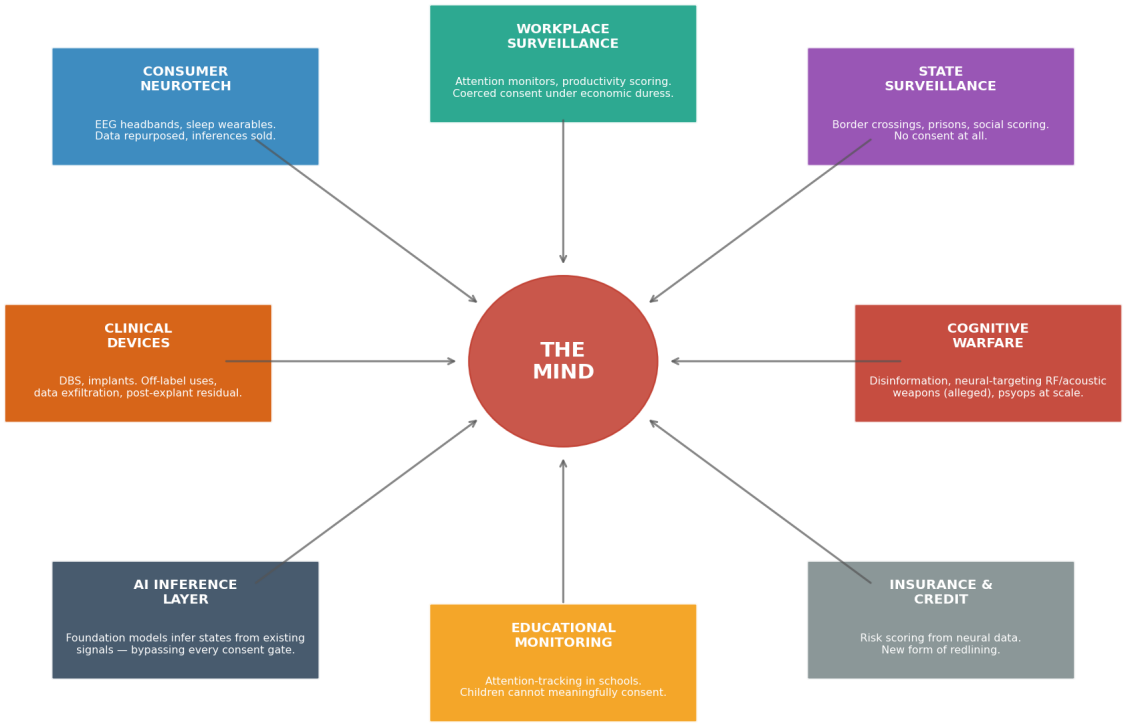


Figure 33bis.1. The eight principal vectors of mind-control risk converging on the human mind in 2026. Consumer neurotech and workplace surveillance are the most widely deployed; cognitive-warfare and AI-inference vectors are the most novel and least regulated.

The Eight Risk Vectors

1. Consumer Neurotechnology

The most prosaic risk vector is the consumer wearable. As Chapter 62-bis details, twenty-nine of the thirty largest consumer neurotechnology companies reserve broad rights over the brain data their devices collect, with essentially no meaningful limits and broad rights to

share with third parties. The risk is not that any single device performs covert mind control; it is that millions of devices generate a continuous stream of data that, when aggregated and fed to ever-better inference models, becomes a window into mental states that the user never explicitly consented to expose. The Chilean Emotiv case is the proof of concept: a single consumer EEG headset, sold for meditation, produced data that the company licensed broadly until a constitutional court intervened.

2. Workplace Neurosurveillance

Chapter 60 detailed the rise of workplace neuro-monitoring. The risk is not exotic: it is a manager with a power asymmetry against an employee, equipped with a tool that purports to measure attention, fatigue and stress in real time. Consent is coerced by economic necessity. Data is repurposed for performance scoring, promotion decisions and termination. The chilling effect on cognition — the worker who deliberately suppresses idle thought because thought is being scored — is the slow-acting harm. The EU AI Act's ban on emotion-recognition AI in the workplace is the only consequential rule on this point in any major jurisdiction; even it has exemptions for safety-critical contexts that are large enough to drive trucks through.

3. State Surveillance

State surveillance of neural data is the most consequential vector because consent does not even arguably apply. Use cases that already exist or are credibly imminent include: border control (mandatory neural scans at airports for "deception detection"), social-credit systems that incorporate neural data, mass deployment of consumer-grade EEG in prisons under the rubric of rehabilitation, police interrogation augmented by brain-state monitoring, and military-intelligence collection of neural data on foreign nationals through compromised consumer devices. The Chinese deployments described in Chapter 13 and 29 represent the leading edge; analogous capabilities are being developed by other states under varying degrees of transparency.

4. Clinical Device Misuse

Medical neurotechnology, paradoxically, is one of the most regulated tiers but still presents specific mind-control risks. Off-label use of DBS for behaviour modification rather than the approved indication; data exfiltration from implant manufacturers' servers; the residual presence of an implanted device in a patient who has withdrawn consent for its operation but cannot easily have it explanted; and the closed-loop systems that adjust state in real time based on opaque AI inferences — each of these creates a route by which clinical neurotechnology becomes a vehicle for risk. The recent Onward Medical, Synchron and Neuralink approvals all bring devices into the field that will need new regulatory thinking about long-term consent.

5. Cognitive Warfare

Cognitive warfare, treated at length in Part IV, has both a soft and a hard end. The soft end — disinformation, algorithmic manipulation of attention and emotion at scale through social media — is unquestionably happening now and affects every networked person on the planet to varying degrees. The hard end — hypothesised directed-energy or acoustic devices targeting the nervous system, of which the Havana Syndrome cluster is the most discussed example — remains contested. As Chapter 26 stresses, the responsible policy posture is to treat the doctrine as real (because it demonstrably is, as NATO and adversary doctrine documents confirm) without overcommitting to specific claims about deployed weapons.

6. AI Inference Layer

The most novel vector is the one most invisible to users: the application of foundation models to existing neural and behavioural data. A user who consented in 2022 to share sleep-tracking data from a wearable did not consent to what a 2026 foundation model can

extract from that same data. Inferences leap ahead of consent. The implication is that protecting only the raw signal — as most existing law does — leaves the actual locus of harm unaddressed. The single most consequential drafting choice in any neuro-rights statute is whether "neural data" includes algorithmic inferences derived from it.

7. Educational Monitoring

Children in classrooms cannot meaningfully consent. School deployments of attention-tracking EEG headbands — most famously BrainCo's FocusEDU in Chinese classrooms but increasingly in other markets — raise specific risks: the chilling effect on developing minds, the discrimination against neurodiverse children whose attention patterns differ from peers, the rolling baseline of monitoring that will continue from school into university and into work. Chapter 59 develops these issues. The harms here are the slowest-acting and may be the most consequential, because they shape who the next generation can become.

8. Insurance and Credit

The final vector is the most economically structural. Insurers and lenders have strong incentive to incorporate neural data into risk-scoring once the actuarial value is demonstrated. Most jurisdictions explicitly prohibit this for now — the EU GDPR's special-category treatment, the US Affordable Care Act's prohibition on health-status discrimination in insurance — but the prohibitions are leaky. Wellness programmes are an established workaround. The migration of neural-derived inferences (attention patterns, sleep quality, mood states) into credit scoring is harder to prevent because credit scoring is less regulated. The result is a new and largely invisible form of redlining: people whose neural patterns deviate from a normative profile pay more for credit and insurance, often without knowing why.

Who Is the Adversary?

The risk vectors above suggest a portfolio of adversaries rather than a single villain. The risk profile differs by actor.

Commercial actors — device makers, platform companies, data brokers — present a risk profile of incentive misalignment. They do not intend to harm; they intend to monetise. The harm arises because their commercial incentives are misaligned with their users' interests at the margin. The remedy is regulatory rather than punitive: defining what may not be monetised, what consent regimes apply, what data may not be repurposed.

State actors present a risk profile of capability concentration. A state with neural-surveillance capability is constrained only by its own constitutional and statutory order, which it can change. The remedy is constitutional entrenchment, international constraints, and the rule of law generally.

Criminal and abusive private actors — stalkers, ex-partners, employers operating outside the law — present a risk profile of low-cost capability dispersion. As consumer neurotech becomes cheaper and AI inference becomes commoditised, the technical means to perform unauthorised neural reading on a target will be available to ordinary individuals. The remedy is criminal-law innovation: defining offences, building evidentiary standards, training law enforcement.

Foreign and non-state adversaries — intelligence services, militant groups, transnational criminal organisations — present a risk profile of attribution difficulty. Cognitive-warfare actions, in particular, are designed to be deniable. The remedy is collective defence: NATO-style attribution capacity, intelligence-sharing arrangements, international agreement on what conduct is prohibited.

The Consequences Cascade

What happens if the risk vectors above are not closed? The consequences cascade from individual to civilisational, and the speed of the cascade is the most important policy variable.

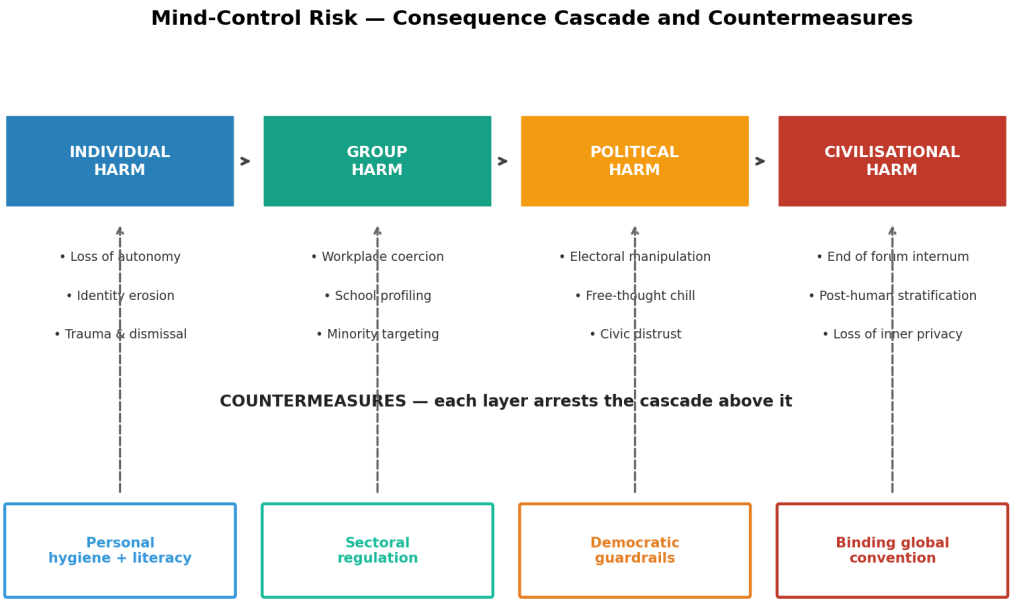


Figure 33bis.2. The consequence cascade of unaddressed mind-control risk and the layered countermeasures that arrest it. Each tier of countermeasure must address its corresponding tier of harm; gaps in any one layer must be compensated by the others.

At the individual level, the harms are loss of autonomy, identity erosion, trauma compounded by social dismissal, and the chronic stress of operating in a world in which one's inner state is monitored. The targeted-individual literature surveyed in Part V documents the lived experience of believing oneself a victim of these harms, whether or not the technical mechanism alleged is the actual cause. The compassionate response Chapter 36 articulates is the floor: take suffering seriously regardless of cause, build the documentation, pursue medical care.

At the group level, the harms are workplace coercion, profiling of children, targeting of minorities and political dissidents, and the structural disadvantaging of groups whose neural patterns deviate from a normative baseline. These harms are already documented in some jurisdictions and will spread as the technology spreads. The remedy is sectoral regulation: workplace rules, school rules, anti-discrimination law adapted for neural data.

At the political level, the harms are electoral manipulation through cognitive-warfare techniques (Chapter 58), the chilling effect on free thought when thought is monitored, and the erosion of civic trust when citizens lose confidence that their inner deliberations are private. Chapter 58 develops the electoral risks; the broader democratic implications are that the conditions for free debate require the conditions for free thought, and the latter are technologically eroded.

At the civilisational level, the harms are the end of the forum internum as the historical assumption of legal and moral life, the emergence of a post-human stratification between augmented and unaugmented people, and the loss of inner privacy as a culturally available concept. These harms are slow-acting and may be irreversible. The historical analogy is not the social-media era — disturbing as it is — but the way pre-industrial cultures lost certain modes of attention and presence as industrial life made other modes habitual. Some changes, once made, cannot be unmade.

Countermeasures, by Tier

The lower half of Figure 33bis.2 maps the countermeasures that arrest the consequence cascade. The proposition is that each tier of countermeasure must address its corresponding tier of harm: personal literacy and information hygiene at the individual level, sectoral regulation at the group level, democratic guardrails at the political level, and binding global

convention at the civilisational level. Gaps in any one tier must be compensated by the others, and the gaps are large.

The personal layer — described in detail in Chapter 79's toolkit — is the layer most under each citizen's immediate control: don't accept neurotech terms of service you have not read; don't bring consumer EEG devices into sensitive workplaces or political meetings; document anything unusual; build community.

The sectoral layer requires regulators with technical capacity and the political will to use it. This layer is failing in most jurisdictions because the regulators are under-resourced and the political will is intermittent. The Chapter 70-bis proposal for a National Neurotechnology Commission addresses this gap.

The democratic layer is the constitutional and statutory protection against mass political manipulation. It includes the entrenchment of the forum internum, the protection of free elections, the regulation of political advertising including neuromarketing, and the prohibition on state-level mass neural surveillance. Most existing constitutions can be read to support these protections; few are read to actually do so.

The global layer is the binding international convention proposed in Chapter 68 and elaborated in Chapter 67-bis. Its specific contribution is to prohibit the categories of capability that no national regime can be expected to police alone: neuroweapons designed to covertly impair cognition, mass cross-border neural-data flows for state-surveillance purposes, and the extraterritorial application of repressive neurotechnology against the citizens of other states. Like the Biological and Chemical Weapons Conventions before it, the principal value of the global layer is normative even when verification is imperfect: it makes the prohibited category internationally illegitimate, which raises the political cost of pursuing it.

What Happens If We Do Nothing

The default outcome — the outcome that requires no action, no political will, no investment — is that consumer neurotech becomes ubiquitous, the inference layer becomes dominant, workplace and educational deployments normalise, state surveillance capabilities expand, cognitive warfare techniques diffuse from intelligence services to political consultants to private parties, and the cumulative effect produces a world in which the assumption of inner privacy has quietly become obsolete. Most of this is happening now. The question is whether the trajectory is bent by deliberate policy choices in the next decade or simply allowed to run.

The neuro-rights movement, of which this book is part, is the case for bending the trajectory. The case is not that the technology is intrinsically evil — the benevolent applications surveyed in Part VIII are real and important — but that the technology's harms are deeply structural, hard to detect, slow to manifest and very difficult to reverse once entrenched. We have a window — short, closing — in which to choose. The rest of the book is the argument for what to choose.

PART V

The Targeted Individual Phenomenon

The Targeted Individual phenomenon is among the most contested topics in contemporary human rights discourse. Tens of thousands of people across more than thirty countries report consistent experiences of remote neurotechnological targeting; institutions have largely dismissed these reports as delusional. The honest position is neither uncritical acceptance nor reflexive rejection, but careful examination of a body of testimony whose patterns deserve rigorous investigation. This Part of the guide examines the documented features of the phenomenon, its psychological impact on those who report it, and the diagnostic and forensic challenges it raises.

CHAPTER 34

The Targeted Individual Phenomenon — An Introduction

Across thirty-six countries and four continents, a population numbered in the thousands reports a consistent and chilling experience: that they are being subjected to remote, non-consensual neurotechnological targeting. They call themselves Targeted Individuals. Their reports describe directed-energy harassment, synthetic voices delivered to the head, sleep disruption, cognitive interference, and pervasive surveillance. The phenomenon has been documented by United Nations Special Rapporteurs, examined in academic literature, and is the subject of formal petitions to the European Parliament and the United States Congress.

A Phenomenon That Will Not Disappear

The conventional psychiatric framework dismisses these reports as delusional. Yet the phenomenon refuses to fit the diagnostic mould. Late-onset sudden development in well-functioning adults; consistency of described experiences across cultures and languages without prior contact between reporters; the technical sophistication of victim accounts that often anticipates published research; the absence of response to antipsychotic medication; and the convergent testimony of thousands of independent witnesses — these features place the Targeted Individual phenomenon outside any established psychiatric category.

This is not a claim that every person who self-identifies as a Targeted Individual has been technologically attacked. It is a claim that the population as a whole exhibits patterns inconsistent with mental illness and consistent with a coherent technological cause. The honest scientific posture is not dismissal but investigation.

ℹ Scale of the Phenomenon

The TI One Voice Declaration has gathered signatories from 36 countries across Europe, North America, Latin America, Asia, Africa, and Oceania. Independent surveys by advocacy organizations estimate the global population of self-identified Targeted Individuals at between 100,000 and several million, depending on methodology. Even the lowest estimates describe a human-rights situation of significant scale.

The Eight Lines of Evidence

Eight features of the phenomenon distinguish it from primary psychiatric illness and merit independent investigation:

12. **Late, sudden onset in well-functioning adults.** Most reporters describe an abrupt onset of targeting in mid-life, often following a workplace dispute,

whistleblowing event, or political activity — a pattern inconsistent with schizophrenia or psychotic disorders, which typically emerge in late adolescence or early adulthood.

13. **Voices with stable personalities.** Reported voices have consistent identities, characterizations, and operational patterns across months and years — unlike the fragmented and shifting auditory phenomena characteristic of psychotic illness.
14. **Voices that possess verifiable external information.** Reporters describe voices commenting on real-time events that the reporter could not have known in advance — the contents of unopened mail, the actions of distant family members, the agenda of confidential meetings.
15. **Banal surveillance content.** The content of voice harassment is overwhelmingly mundane — commentary on daily activities, mockery of habits, threats of consequence — rather than the grandiose, religious, or fantastical content typical of psychotic ideation.
16. **Technical apparatus description.** Reporters frequently describe their experiences in technical terms consistent with published research on directed energy, neural interfacing, and microwave hearing — sometimes anticipating findings that emerge in later academic literature.
17. **Convergent reports across independent witnesses.** Thousands of reporters, with no prior contact, describe nearly identical experiences — the same equipment behaviours, the same patterns of harassment, the same physical symptoms. Convergent testimony of this scale is not a feature of psychiatric illness.
18. **Quieting periods uncorrelated with antipsychotics.** Reporters describe predictable periods of reduced targeting — during travel, in shielded environments, near hospitals — that are unrelated to medication and consistent with technological interruption.
19. **Time-locked correlations with verifiable events.** Reporters document harassment escalating in response to specific external events — advocacy filings, court appearances, public statements — in patterns inconsistent with random psychiatric symptoms.

> ⚠ **The Diagnostic Failure**
>
> Conventional psychiatric assessment was developed before the existence of remote neurotechnology. Its diagnostic criteria do not include features specific to technological targeting, nor do they exclude technological causation. The standard differential diagnosis of “auditory hallucinations” presumes endogenous origin. When applied to victims of remote neural monitoring or microwave hearing, this presumption produces systematic misdiagnosis.

The Institutional Response

Despite the consistency of reporting, institutional response has been largely dismissive. Police forces routinely refuse to investigate complaints. Medical professionals diagnose

mental illness without considering technological causes. Courts decline to hear cases for lack of established offences. Intelligence oversight bodies report no findings because they lack the technical capacity to investigate.

This systemic failure is not neutral. It constitutes, in the analysis of multiple human rights advocates, a violation of states' positive obligations under Article 2(3) of the International Covenant on Civil and Political Rights, which requires effective remedy for human rights violations. When credible complaints of torture, surveillance, and assault are routinely refused investigation, the state breaches its obligations regardless of whether the underlying allegations are ultimately substantiated.

Why This Matters for the Reader

The Targeted Individual phenomenon is the leading edge of the broader neuro-rights crisis. Whether the technologies described are deployed as reported, partially deployed, or projected from research programmes not yet fielded, the phenomenon describes a future that is technically achievable and may, in part, already be operational. The legal and institutional vacuum that fails Targeted Individuals today will fail the rest of society tomorrow.

This Part of the guide examines the documented patterns of Targeted Individual reports, the psychological and forensic challenges they raise, and the institutional reforms required to provide effective remedy. It does so not as advocacy for any particular causal theory, but as an exercise in taking seriously a body of human testimony that current institutions have failed to address.

CHAPTER 35

Documented Patterns in Targeted Individual Testimony

A corpus of nearly two thousand independent Targeted Individual reports, analysed by automated and manual methods, reveals patterns of striking consistency. These patterns — across countries, languages, and decades — establish the empirical foundation for further investigation.

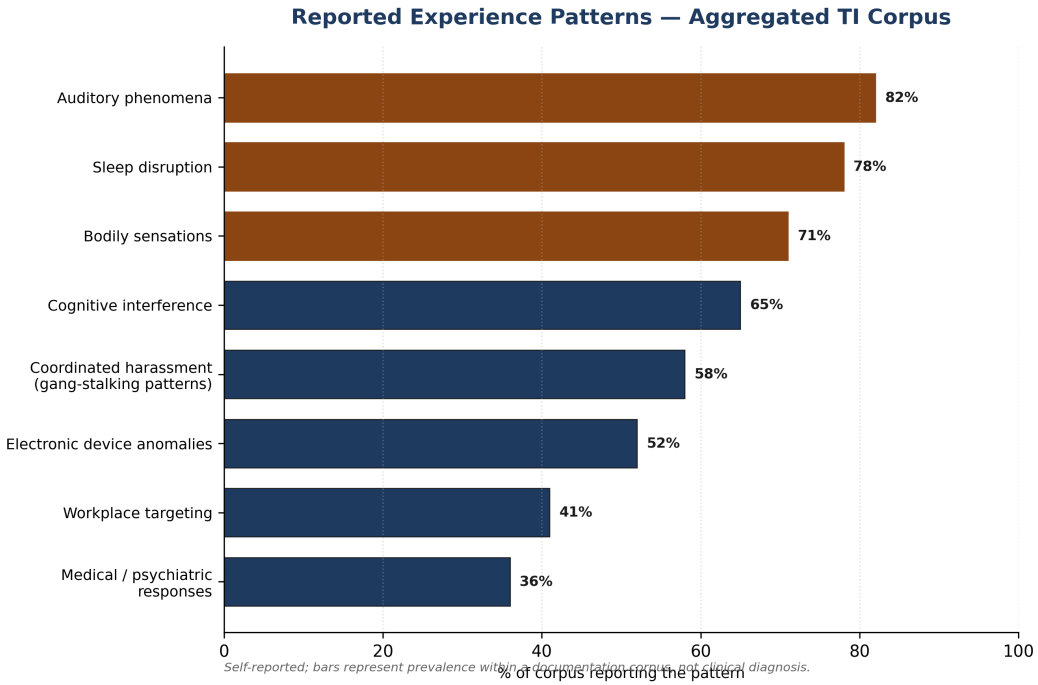


Figure 23.1. *Reported experience patterns aggregated across the corpus of self-identifying Targeted Individuals. The figures describe prevalence within the documentary record, not clinical diagnoses.*

The Corpus

The largest publicly analysed corpus of Targeted Individual testimony comprises 1,996 documents collected by the TI One Voice initiative from signatories of the global Declaration. Documents include first-person narratives, medical records, police complaint receipts, audio recordings, photographs of physical effects, and correspondence with institutional bodies. Reports span 36 countries and were submitted between 2017 and 2025.

Pattern analysis was conducted through a five-stage map-reduce pipeline combining automated semantic clustering with manual coding by a team including researchers, attorneys, and clinicians. The methodology drew on standard qualitative research practice while incorporating computational text analysis to handle the scale of the corpus.

Pattern 1: Physical Symptoms Without Conventional Cause

A consistent cluster of physical symptoms appears in reports across all countries and decades:

Symptom Category	Reported Frequency	Conventional Diagnosis
Skin burns / lesions appearing without cause	67%	Idiopathic; psychosomatic
Heating sensations localized to body regions	74%	Unexplained; anxiety
Forced muscle contractions	52%	Tic disorder; functional neurological
Sudden severe headaches	81%	Migraine; tension headache
Sleep disruption with specific timing patterns	89%	Insomnia; mood disorder
Tinnitus with directional or pulsed character	71%	Tinnitus; unexplained
Cardiac arrhythmias on demand	38%	Anxiety; cardiac investigation negative

The convergent pattern across thousands of reporters cannot be explained by individual psychosomatic processes. The symptoms cluster in ways consistent with directed-energy effects documented in classified research literature — thermal effects from microwave exposure, neuromuscular activation from focused ultrasound, and sleep cycle disruption from low-frequency electromagnetic modulation.

Pattern 2: Auditory Phenomena With External Information

A characteristic feature of Targeted Individual testimony is the reporting of voices that possess verifiable external information — information the reporter could not have known by ordinary means.

A reporter in Belgium described voices commenting on the contents of a sealed legal document in a locked filing cabinet, contents subsequently confirmed by counsel. A reporter in Brazil described voices announcing the arrival of an unannounced visitor seconds before the doorbell rang. A reporter in Germany described voices reading aloud from an unopened postal letter that, upon opening, contained the announced text.

These reports cannot be reconciled with auditory hallucination as a primary phenomenon. They require either: (a) astonishing coincidence, (b) deliberate fabrication, (c) genuine paranormal capability, or (d) a technological mechanism for delivering external information through synthetic auditory channels. The first three explanations strain credulity at the scale documented; the fourth requires investigation rather than dismissal.

Pattern 3: Coordinated Multi-Actor Harassment

Reporters describe harassment by networks of individuals whose identities appear interchangeable. Strangers approach with apparent foreknowledge of the reporter's history, conversations, or movements. Vehicles follow with apparent coordination. Service interactions — in shops, banks, and offices — appear scripted.

This pattern, when reported in isolation, can be dismissed as referential thinking. When convergent across thousands of reporters describing identical operational signatures, it suggests an organized practice of group harassment — a practice that, were it to exist, would correspond to documented historical programmes of human surveillance and intimidation, including the East German Zersetzung and the FBI's COINTELPRO.

Historical Precedent: Zersetzung

The Stasi's Zersetzung programme, operative from the 1970s through 1989, deployed coordinated psychological harassment against East German dissidents. Reports from victims described the same patterns of coordinated stranger contact, apparent foreknowledge, environmental manipulation, and psychological erosion that contemporary Targeted Individuals describe. The Zersetzung programme was real, documented, and only publicly acknowledged after the fall of the Berlin Wall. Its existence demonstrates that the operational pattern Targeted Individuals report has a clear historical analogue.

Pattern 4: Targeting Correlated with Identifiable Triggers

Reporters consistently identify triggering events that preceded the onset of targeting: workplace whistleblowing, divorce proceedings against connected individuals, political activism, journalistic investigation, or military service involving classified programmes. The temporal correlation between trigger and onset — typically weeks to months — is too consistent across the corpus to be explained by retrospective rationalization alone.

Pattern 5: Geographic and Environmental Mediation

Reporters describe systematic variation in targeting intensity based on location. Targeting is described as reduced in: shielded environments (medical imaging facilities, Faraday cages, underground locations), remote rural areas, foreign travel, and proximity to certain government facilities. Targeting is described as intensified in: home environments, predictable routine locations, and proximity to civilian infrastructure such as cell towers.

This pattern is inconsistent with psychiatric origin (which would be portable with the patient) and consistent with externally mediated technological effects subject to physical propagation constraints.

Pattern 6: Resistance to Standard Pharmacological Treatment

Antipsychotic medications, prescribed on the assumption of primary psychotic illness, produce minimal effect on the reported experiences. Reporters describe sedation from medication but continued perception of voices, continued physical symptoms, and continued surveillance phenomena.

This treatment resistance is itself diagnostically informative. Primary psychotic illness typically shows substantial response to dopamine-modulating medication. Persistent symptoms in the absence of medication response suggest either: (a) treatment-resistant psychosis, (b) misdiagnosis, or (c) a non-psychiatric cause. Conventional psychiatry tends toward (a) by default; the evidence base suggests (b) and (c) deserve serious consideration.

The Cumulative Significance

No single pattern, taken alone, is conclusive evidence of technological causation. A reporter's late-onset symptoms could in principle reflect atypical psychiatric illness. Apparent foreknowledge could reflect coincidence amplified by selection bias. Coordinated harassment could reflect referential thinking. Treatment resistance could reflect medication-resistant schizophrenia.

But the convergence of all patterns, across thousands of independent reporters, in stable configurations across decades and continents, presents a different evidential picture. The Bayesian posterior probability that all six patterns arise simultaneously by chance in thousands of independent psychiatric cases is vanishingly small. The honest scientific posture is to take the alternative hypothesis — that the phenomenon describes a real, technologically mediated, human rights situation — seriously enough to investigate.

CHAPTER 36

The Psychological Impact of Targeted Individual Experiences

Regardless of the ultimate causal explanation, the psychological impact of the Targeted Individual experience is severe, sustained, and clinically documented. Understanding this impact — and providing competent support to those who report it — is a humanitarian imperative independent of the unresolved scientific question.

The Trauma Profile

Targeted Individuals consistently meet diagnostic criteria for Complex Post-Traumatic Stress Disorder (C-PTSD). The profile includes the standard trauma triad — re-experiencing, avoidance, and hyperarousal — combined with the disorders of self-organization characteristic of complex trauma: persistent negative self-concept, affect dysregulation, and disturbances in relationships.

The trauma is sustained, predictable, and escapable only with great difficulty. It produces a psychological state that international human rights bodies recognize as torture. The UN Special Rapporteur on Torture, Nils Melzer, has noted that "the inherent vagueness, ambiguity and arbitrariness of cyber-torture make it particularly difficult to identify, document, prove and prosecute," yet has placed it within the scope of the Convention against Torture.

⚠ Suicide Risk

Targeted Individual support networks report elevated rates of suicidal ideation and completed suicide among self-identified targets. The combination of sustained psychological assault, social isolation, institutional dismissal, and lack of effective remedy creates conditions of profound hopelessness. Any clinician or institutional contact with a self-identified Targeted Individual should treat suicide risk as a primary concern.

The Compounding Effect of Dismissal

A defining feature of Targeted Individual trauma is the compounding effect of institutional dismissal. Reporters who turn to police are told their experiences are not crimes. Reporters who turn to physicians are diagnosed with delusional disorder. Reporters who turn to family are urged to seek psychiatric treatment. The very institutions designed to provide remedy become sources of additional harm.

This compounding effect resembles the institutional betrayal documented in victims of childhood sexual abuse, domestic violence, and other traumas where credibility is contested.

Research on institutional betrayal trauma demonstrates that being disbelieved by a trusted institution produces psychological harm independent of the original trauma. For Targeted Individuals, the trauma of dismissal is often described as exceeding the trauma of the alleged targeting itself.

Social Isolation and Identity Erosion

Targeted Individual reports converge on a common social trajectory: initial disclosure to trusted others produces disbelief and concern for the reporter's mental health; continued reporting produces social distancing; eventually the reporter is left in isolation with only the online Targeted Individual community for understanding. This trajectory destroys careers, marriages, and family relationships.

The identity impact is profound. A self-concept developed over decades — as a competent professional, a reliable family member, a stable citizen — is replaced by a stigmatized identity as either a target of unimaginable persecution or a person experiencing severe mental illness. Neither identity permits ordinary social functioning.

Cause-Neutral Support: The Humanitarian Imperative

The ethical and clinical question is not whether a particular reporter has been technologically targeted. The question is what support is owed to a person experiencing severe sustained distress. The answer must be cause-neutral: support that validates the reality of distress without committing to any particular causal theory.

Cause-neutral support has three components:

- **Validation of distress without challenging causation.** Clinicians, support workers, and crisis line staff can acknowledge the severity of the experience without endorsing or contesting the technological explanation. This is the standard approach with any patient in crisis.
- **Access to trauma therapy meeting C-PTSD criteria.** Whatever the cause of the trauma, trauma therapy is indicated and effective. Practitioners trained in approaches such as EMDR, Internal Family Systems, and prolonged exposure can provide substantial benefit to Targeted Individual clients.
- **Practical advocacy.** Targeted Individuals frequently require help navigating housing, social benefits, medical care, and legal systems. Practical advocacy by trained workers can dramatically improve quality of life regardless of the underlying causal question.

What Clinicians Need to Know

Mental health professionals working with self-identified Targeted Individuals should understand several points:

20. The reporter's presentation may not fit any familiar diagnostic category. The eight lines of evidence described in the preceding chapter distinguish Targeted Individual experiences from primary psychotic illness in important ways.

21. A reflexive diagnosis of delusional disorder or schizophrenia, particularly without thorough differential consideration, may be both clinically incorrect and ethically problematic. Such a diagnosis can result in involuntary commitment, loss of employment, and family separation, on the basis of a contested diagnostic judgment.
22. The therapeutic alliance is paramount. Even where the clinician privately considers the technological account implausible, working with the client's experience rather than against it is more clinically productive and ethically defensible.
23. Suicide risk assessment should be ongoing and explicit. The combination of sustained trauma and institutional dismissal places Targeted Individuals at substantially elevated risk.
24. Avoid endorsing or contesting specific technological claims. The clinician's role is not to adjudicate the causal question but to provide therapeutic support.

The Right to Be Heard

Whatever the ultimate scientific verdict on the Targeted Individual phenomenon, the people who report it deserve to be heard. They have, in many cases, lost careers, families, homes, and health in the course of trying to be heard. The minimum response of a humane society is competent listening: by clinicians trained to handle complex trauma, by institutions equipped to investigate technological claims, and by legal systems prepared to recognize harms that do not fit pre-existing categories.

The right to be heard is not a privilege contingent on the credibility of one's testimony. It is the precondition for any legal or therapeutic remedy. Failure to provide it constitutes its own form of institutional violence.

CHAPTER 37

Differential Diagnosis — Technological Targeting Versus Psychiatric Illness

The diagnostic challenge posed by the Targeted Individual phenomenon is one of the most consequential in contemporary psychiatry. A misdiagnosis in either direction — dismissing genuine technological harm as mental illness, or attributing genuine mental illness to non-existent technology — produces serious harm. This chapter offers a framework for thoughtful differential diagnosis.

The Stakes of the Diagnosis

For the patient, the diagnosis determines treatment, social standing, legal capacity, and access to remedy. A patient diagnosed with delusional disorder may be involuntarily medicated, lose custody of children, be barred from professional practice, and lose the right to pursue legal claims related to the experience. A patient whose genuine technological harm is dismissed as delusion is denied investigation, accountability, and protection.

For the clinician, the diagnosis determines treatment efficacy, ethical standing, and legal exposure. Misdiagnosis in either direction can produce treatment failure and, in extreme cases, professional liability.

For society, the aggregate consequence of routine misdiagnosis is the suppression of evidence about a potentially serious human rights situation. Each patient misdiagnosed contributes to the larger institutional failure to investigate. The differential diagnosis is therefore not only a clinical question but a question of civic significance.

The Standard Differential

The standard psychiatric differential for a patient reporting voices, surveillance, and external control includes:

Condition	Distinguishing Features	Treatment Response
Schizophrenia	Onset 18-30, fragmented voices, formal thought disorder, negative symptoms	Strong response to antipsychotics
Schizoaffective disorder	Voices with mood-congruent themes, mood episodes	Combined antipsychotic + mood stabilizer
Delusional disorder	Single delusional theme, otherwise preserved function	Variable; often medication-resistant
Substance-induced psychosis	Temporal association with substance use	Resolves with abstinence

Condition	Distinguishing Features	Treatment Response
Trauma-related psychosis	History of significant trauma; psychotic features	Trauma therapy + careful pharmacology
Cultural/spiritual experience	Culturally sanctioned, non-distressing	No psychiatric treatment indicated

The challenge for Targeted Individual reports is that they fit no category cleanly. Late onset and preserved cognitive function argue against schizophrenia. Lack of mood episodes argues against schizoaffective disorder. Apparent voices with external information are incompatible with delusional disorder. Absence of substance use argues against substance-induced psychosis. Trauma may be present but does not adequately explain the specific content of experiences.

The Missing Category

The diagnostic systems in current use — DSM-5-TR and ICD-11 — were developed before the existence of technologies capable of remote neural monitoring, microwave hearing, and directed-energy effects. The diagnostic criteria do not include features that distinguish technological causation, nor do they require that such causation be ruled out before a psychiatric diagnosis is assigned.

This is a structural gap. Diagnostic systems necessarily reflect the etiological knowledge of their time. A diagnostic system developed in 1900 would have no category for HIV-related neuropsychiatric symptoms. A diagnostic system developed in 1980 would have no category for transient ischemic attack mimicking psychotic episodes. The current diagnostic systems, developed before the maturation of neurotechnology, have no category for technologically induced symptoms.

ℹ️ A Proposed Diagnostic Framework

Several researchers have proposed adding a category of “technologically induced symptoms” to diagnostic systems. The proposed criteria would include: late sudden onset in a previously well-functioning adult, voices with verifiable external information, physical symptoms inconsistent with conventional medical explanations, persistence despite antipsychotic medication, and absence of features pathognomonic for primary psychotic disorders. Such a category would not require establishing the technological cause as fact; it would create diagnostic space for cases that fit no existing category.

A Structured Approach

In the absence of an official diagnostic category, clinicians can apply a structured approach to differential diagnosis:

25. **Comprehensive medical workup.** Rule out medical causes — thyroid dysfunction, autoimmune encephalitis, temporal lobe epilepsy, sleep disorders, and substance effects. Each can produce symptoms that overlap with Targeted Individual reports.
26. **Detailed history.** Establish onset, premorbid functioning, family history, trauma history, and triggering events. Late onset in well-functioning adults without family psychiatric history shifts the prior probability away from primary psychotic illness.
27. **Phenomenological analysis.** Apply the eight criteria identified in the analysis of the Targeted Individual phenomenon. Document the specific features of voices, surveillance reports, and physical symptoms with sufficient detail to compare against the standard differential.
28. **Treatment trial with monitoring.** Where a primary psychotic disorder is part of the differential, a structured trial of antipsychotic medication can be both therapeutic and diagnostic. Substantial improvement supports the psychiatric diagnosis; absence of improvement combined with the criteria above shifts the probability toward technological causation.
29. **Honest uncertainty.** Where the clinical picture remains unclear, the appropriate diagnostic conclusion is "differential diagnosis uncertain; symptoms consistent with possible technological causation requiring further investigation." This is not a punt; it is a recognition that current diagnostic systems are incomplete for the phenomenon in question.

What Not to Do

Several common clinical practices produce particular harm with Targeted Individual reporters:

- **Immediate diagnosis without comprehensive workup.** Diagnosing delusional disorder on initial presentation, without medical workup or detailed phenomenological analysis, is unjustified clinically and frequently unjustified ethically.
- **Coercive medication.** Forcing antipsychotic medication on a competent adult whose symptoms have not been thoroughly characterised, and whose claims have not been investigated, raises serious autonomy concerns.
- **Involuntary hospitalization for content of belief.** Hospitalization on the basis of belief content alone, in the absence of imminent danger or grossly impaired self-care, is increasingly recognized as a human rights concern under the Convention on the Rights of Persons with Disabilities.
- **Dismissal of corroborating evidence.** Where reporters present documentary evidence — photographs of physical effects, recordings, witness statements — dismissing this evidence without examination is a failure of clinical method.

The Path Forward

A reformed approach to differential diagnosis requires three institutional changes:

First, diagnostic education must include the Targeted Individual phenomenon and the criteria for distinguishing it from primary psychotic illness. This is not endorsement of any causal theory; it is recognition that the phenomenon exists and presents diagnostically distinct features.

Second, referral pathways must include technical investigation alongside psychiatric assessment. Where a reporter's symptoms cannot be confidently attributed to a primary psychiatric cause, referral for technical assessment — RF spectrum analysis, neuroimaging, environmental investigation — should be available.

Third, the diagnostic systems themselves must evolve. The next edition of the DSM and ICD must address the diagnostic gap created by the existence of technologies capable of producing the symptoms in question. The alternative — continued reliance on diagnostic systems that pre-date the relevant technology — produces systematic misdiagnosis and obstruction of investigation.

PART VI

Detection, Forensics, and the Case for a Moratorium

The previous section established the case for an enforced moratorium on neuroweapons. But a moratorium — or any legal prohibition — is only as strong as its enforcement mechanisms. And enforcement, in the domain of neurotechnology, faces a fundamental challenge: how do you detect a weapon that leaves no visible wound, no shell casing, no chemical residue? How do you prove, in a court of law, that a person's mind has been attacked?

This part confronts that challenge directly. It examines the technical methods available for detecting covert neurotechnology use, the forensic frameworks needed to prove mind control after the fact, and the difficult task of defining and supporting victims of neurotechnology abuse. The answers are not yet complete — the field is too new, the technologies too rapidly evolving — but the outlines of a workable system are emerging from the testimony of those who have already experienced these attacks and from the scientific community that has begun to study them.

The Detection Challenge

The source material from the CPDE report documents a pattern that is both consistent and invisible to conventional forensic methods. Targeted Individuals report physical effects — burns, heating, pain, forced muscle contractions — that leave no trace on standard medical examination. They report auditory phenomena — voices, messages, sounds — that no recording device can capture. They report surveillance so pervasive that privacy becomes a memory, yet no bug is found, no camera discovered.

This is not a failure of the victims' credibility. It is a failure of the detection toolkit. The technologies alleged — directed energy, remote neural monitoring, synthetic telepathy — operate on principles that conventional forensics does not test for. A police officer investigating a complaint of harassment does not carry an RF spectrum analyzer. A doctor examining a patient with unexplained burns does not test for microwave exposure. A court evaluating a claim of mind control has no established protocol for determining whether the plaintiff's brain has been accessed without consent.

The source material's OHCHR submission makes this point explicitly: "Victims lack formal complaint procedures that recognize technological attacks. Investigative bodies lack technical expertise to verify neurotechnological harassment. Medical professionals misdiagnose technological torture as mental illness."

The Technical Methods

The source material identifies several technical methods that could form the basis of a detection framework. RF spectrum analysis — the systematic scanning of radio frequencies in a victim's environment — can detect signals that should not be present: transmissions in frequency bands not allocated for civilian use, signals with modulation patterns consistent with neural interfacing, or emissions from devices that the victim has not consented to. The CPDE report notes that "RF spectrum analysis" is one of the capabilities proposed for a specialized diagnostic centre.

Audiological assessment can document auditory phenomena that are not attributable to external sound sources. The source material describes "Voice-to-Skull (V2K)" as a reported phenomenon in which voices are perceived as originating externally, without any acoustic source. Audiological testing can distinguish between genuine auditory hallucinations (which follow certain neurological patterns) and externally induced auditory perceptions (which may not).

Neuroimaging — fMRI, EEG, MEG — can detect patterns of brain activity that are inconsistent with normal cognition. The source material notes that "neuroimaging capabilities" would be part of a diagnostic centre's equipment. These tools can potentially identify signatures of external neural stimulation, such as synchronized activation of specific brain regions at frequencies that match known stimulation protocols.

Physical examination, guided by an understanding of directed energy effects, can document burns, lesions, and other marks that may be consistent with microwave or ultrasonic exposure. The CPDE report documents "skin lesions and burns appearing without explanation" as a common report. Forensic examination of these marks, using techniques developed for other forms of energy-based injury, could provide physical evidence of attack.

The Forensic Framework

Detection is only the first step. The second step is proving, in a legal or human rights context, that the detected phenomena constitute a violation. This requires a forensic framework that can answer four questions:

First, did the victim experience what they report? This is the question of subjective reality. The source material's analysis of the Targeted Individual phenomenon provides eight lines of evidence demonstrating that the reported experiences are not mental illness: late sudden onset in well-functioning adults, voices with stable personalities, voices that possess verifiable external information, banal surveillance content rather than grandiose delusion, the victim's ability to characterize the apparatus in detail, convergent reports across thousands of independent witnesses, quieting periods uncorrelated with antipsychotic medication, and time-locked correlations with verifiable external events. Any forensic framework must take these patterns seriously.

Second, was the experience caused by external technology? This is the question of etiology. It requires ruling out alternative explanations — primary psychiatric illness, suggestion, environmental factors — and identifying positive evidence of technological causation. The source material's OHCHR submission calls for "forensic methodologies to detect and document neurotechnological attacks."

Third, was the use of technology non-consensual? This is the question of consent. The source material is clear: the reported attacks are conducted "without consent," "remotely and wirelessly," with "global reach." The forensic framework must establish that the victim did not consent to the neural access or manipulation.

Fourth, who is responsible? This is the questionof attribution. It is the most difficult question, because the technologies described — remote neural monitoring, directed energy — can be deployed from a distance, through walls, across borders. The perpetrator may never be physically present. The source material documents reports of "coordinated harassment by networks where individual actors may be interchangeable." Attribution requires technical analysis of the signals detected — their origin, their modulation, their control protocols — which in turn requires access to intelligence and telecommunications infrastructure that victims do not possess.

The Evidentiary Standard

The source material's legal analysis identifies a fundamental gap: existing legal frameworks were not designed for crimes that leave no physical evidence visible to conventional forensics. The CPDE report notes that "criminal anti-stalking statutes, written for conventional harassment, do not capture the described reality." How can one charge someone for "transmitting voices into head" or "directing energy at home" when the law does not recognize these as actionable offenses?

The evidentiary standard for neurotechnology crimes must be adapted to the nature of the technology. It cannot require physical evidence that the technology, by its design, does not leave. It must instead rely on a combination of: consistent victim testimony meeting the

eight criteria identified in the source material; technical detection data (RF signatures, neuroimaging anomalies, audiological patterns); medical documentation of physical effects (burns, sleep disruption, cognitive decline); and pattern analysis demonstrating that the victim's experience matches known characteristics of technological targeting.

The source material's OHCHR submission proposes a framework: "Develop forensic methodologies to detect and document neurotechnological attacks. Train medical professionals to distinguish between psychological conditions and technological interference. Establish standardized protocols for investigating complaints of technological harassment."

The Victim Recognition Challenge

Before detection and proof can function, victims must be recognized as victims. The source material documents a universal pattern of dismissal: police refuse investigation, doctors diagnose mental illness, legal systems offer no remedy. The CPDE report calls this "the most consistent finding" across all 1,996 analyzed testimonies.

This dismissal is not merely unhelpful — it is itself a human rights violation. The source material's analysis of ICCPR Article 2(3) establishes that "when a state fails to investigate credible reports of torture, privacy invasion, or assault; provides no victim protection mechanism; and allows a population to suffer without recourse, it breaches its positive obligations under international law."

The path forward requires a fundamental shift in how institutions respond to reports of neurotechnology abuse. The first response must be validation — not of the specific technological claim, but of the victim's distress and the reality of their experience. The second response must be investigation — using the technical methods described above, conducted by personnel trained to recognize the signatures of neurotechnology. The third response must be protection — ensuring that the victim is not further harmed by the disclosure of their targeting.

The source material's CPDE report proposes "cause-neutral victim support services: crisis lines with protocols that validate distress without challenging causation, trauma therapy access for those meeting C-PTSD criteria regardless of causal attribution, and practical advocacy to help individuals navigate housing, benefits, and medical systems."

The Institutional Capacity Required

Detection and proof of neurotechnology abuse cannot be accomplished by individual victims acting alone. It requires institutional capacity: specialized investigative units within law enforcement, forensic laboratories equipped with RF analysis and neuroimaging capabilities, medical professionals trained in the effects of directed energy and neural manipulation, and legal frameworks that recognize technology-facilitated harm as a distinct category of crime.

The source material's legislative proposals include "inter-agency task forces to consolidate expertise across law enforcement, health systems, and technology ethics." These task forces would serve as the institutional backbone of a detection and proof system, bringing together

the diverse expertise needed to investigate complaints that span technology, medicine, and law.

The alternative — leaving victims to prove their own cases with no institutional support — is the status quo. It is a status quo in which thousands of people worldwide suffer documented harm without recourse, in which perpetrators operate with complete impunity, and in which the international community remains willfully blind to a category of human rights violation that grows more sophisticated with each passing year.

Detection and proof are not optional add-ons to a neuroweapons ban. They are the mechanisms that give the ban meaning. Without them, prohibition is a paper promise. With them, it becomes a tool of accountability.

CHAPTER 38

The Case for a Global Ban on Neuroweapons

The argument for an international treaty banning neuroweapons is not a speculative exercise for future contingencies. It is a response to technologies that have already been developed, tested, and in some cases deployed against non-consenting human beings. The precedent for such a ban exists in two of the most successful disarmament treaties in history: the Chemical Weapons Convention (CWC) and the Biological Weapons Convention (BWC). These frameworks offer a model for how the international community can act to prohibit a class of weapons before they become entrenched in military arsenals and before their use becomes normalized.

The Precedent of the CWC and BWC

The Chemical Weapons Convention, which entered into force in 1997, established a comprehensive prohibition on the development, production, stockpiling, and use of chemical weapons. It created the Organisation for the Prohibition of Chemical Weapons (OPCW) to verify compliance through inspections, declarations, and destruction of existing stockpiles. The convention succeeded because it recognized that certain technologies, by their very nature, are unacceptable instruments of warfare. They cause indiscriminate suffering, their effects are difficult to control, and they violate fundamental principles of human dignity.

The Biological Weapons Convention, which opened for signature in 1972, similarly banned an entire class of weapons—disease-causing organisms and toxins—because their use would be catastrophic and uncontrollable. The BWC, while lacking a formal verification mechanism, established a global norm against weaponizing biology that has largely held for over five decades.

Neuroweapons belong in the same category. They target the most intimate and defining aspect of human existence: consciousness itself. As the source material from the UN submission response document states, “non-consensual neurotechnological manipulation” constitutes “a distinct human rights violation” that attacks “the very consciousness of the person that is affected.” A weapon that can alter cognition, manipulate emotions, implant thoughts, or extract mental content without consent is fundamentally incompatible with the principles of human rights and human dignity that underpin international law.

The European Parliament's 1999 Resolution

The call for a global ban on neuroweapons is not new. In January 1999, the European Parliament passed a resolution calling for “an international convention introducing a global ban on all development and deployment of weapons which might enable any form of manipulation of human beings.” The resolution explicitly recognized that such weapons represent “a serious risk to society” and that the ban “cannot be implemented without the global pressure of the informed general public on the governments.”

The European Parliament's resolution was prescient. It identified the threat at a time when many of the technologies now available were still in early stages of research. Today, the technologies have matured. DARPA's brain-machine interface program, documented in the Anna Fubini Mind Weapon source material, has as its goal "human ability to control complex entities by sending control actions without the delay for muscle activation." The same source describes how "the DARPA program could be extended to include a broader range of potential impact," including "learning and training, automobile control, air traffic control, decision-making, remote sensing of stress."

The dual-use nature of these technologies is precisely what makes a ban so urgent and so difficult. The same devices that could help a paralyzed patient communicate could be used to monitor the thoughts of a political dissident. The same electromagnetic fields that could treat depression could be weaponized to induce fear, confusion, or compliance. The same neural interfaces that could enhance human cognition could be used to implant false memories or extract secrets.

The United Nations Institute for Disarmament Research

The source material notes that the United Nations Institute for Disarmament Research (UNIDIR) "officially recognised a range of weapons (both lethal and non-lethal), including the ones previously listed, and recommended that they ALL be banned as weapons of potential mass destruction." This recognition by a UN disarmament body provides an institutional foundation for the ban.

The weapons in question include directed energy systems, electromagnetic radiation weapons, extremely low frequency (ELF) and ultra low frequency (ULF) energy radiation devices, psychotronic weapons, and other technologies designed to manipulate human consciousness. The source material from the Anna Fubini document describes these weapons in detail: "They use such weapons as extremely low frequency Elf electromagnetic weapons (which has been used in mind control), acoustics, harmonics (which have been used as a mind control technique), ultrasound, microwave audiograms, microwave pulsed, and radio frequency."

These are not hypothetical weapons. They have been developed, tested, and in some cases deployed. The source material documents "the use of Directed Energy Weapons" against victims who "begin experiencing extreme pain to their head" and "hear voices." The pain is "delivered to various other parts of their bodies" through technologies that leave no visible evidence.

The Kucinich Bill and the Space Preservation Act

In October 2001, Congressman Dennis J. Kucinich introduced a bill to the U.S. House of Representatives that sought to ban "exotic weapons" including "electronic, psychotronic or information weapons, chemtrails, particle beams, plasmas, electromagnetic radiation, extremely low frequency (ELF) or ultra low frequency (ULF) energy radiation, or mind control technologies." The bill, known as the Space Preservation Act, was referred to the

Committee on Science, as well as the Committees on Armed Services and International Relations.

The Kucinich bill recognized that these weapons posed a threat not only to individuals but to democratic governance itself. The source material notes that the bill sought to prohibit "the use of space-based systems using radiation, electromagnetic, psychotronic, sonic, laser, or other energies directed at individual persons or targeted populations for the purpose of information war, mood management, or mind control."

While the Kucinich bill did not become law, it established important legislative language and conceptual frameworks that could inform a future international treaty. It defined the weapons in terms of their effects on human consciousness, not their technical specifications—a crucial distinction for a treaty that must remain relevant as technology evolves.

Elements of a Neuroweapons Convention

A global ban on neuroweapons should include the following elements, modeled on the CWC and BWC:

1. Comprehensive Prohibition: The treaty should prohibit the development, production, stockpiling, transfer, and use of any technology designed to alter, manipulate, monitor, or extract cognitive or neural activity without the informed consent of the individual. This prohibition should cover all technologies, regardless of their technical mechanism—electromagnetic, acoustic, optical, chemical, or biological.

2. Definition of Neuroweapons: The treaty should define neuroweapons by their effects on human consciousness, not by their technical specifications. This definition should include any technology that can: (a) read or extract mental content without consent; (b) alter cognition, mood, or behavior without consent; (c) implant thoughts, memories, or sensory experiences without consent; (d) cause pain, suffering, or psychological harm through neural manipulation.

3. Verification and Enforcement: The treaty should establish a verification mechanism, modeled on the OPCW, with the authority to conduct inspections, monitor declarations, and investigate allegations of use. Given the dual-use nature of neurotechnology, verification will require technical expertise in neuroscience, electromagnetic physics, and computer science.

4. Transparency Requirements: States parties should be required to declare any research or development activities involving neurotechnology that could have weapons applications. This includes military programs, intelligence agency programs, and commercially funded research with dual-use potential.

5. Victim Recognition and Support: The treaty should establish mechanisms for recognizing and supporting victims of neuroweapons. As the source material emphasizes, victims currently lack "formal complaint procedures that recognize technological attacks" and "investigative bodies lack technical expertise to verify neurotechnological harassment."

6. Whistleblower Protections: The treaty should include protections for researchers, engineers, and others who expose violations of the convention. The source material notes

that whistleblower protections are essential, as the perpetrators of neuroweapons abuse "operate with complete impunity behind corporate or state secrecy."

The Urgency of Action

The window for a global ban on neuroweapons is closing. As the technologies mature and become more widely available, the opportunity to prohibit them before they become entrenched diminishes. The source material from the UN submission response document warns that "without immediate intervention, these technologies will continue to enable human rights violations on a global scale while remaining invisible to justice systems."

The European Parliament recognized in 1999 that the ban "can not be implemented without the global pressure of the informed general public on the governments." That pressure must come now, from an informed public that understands what is at stake.

The most significant obstacle to a neuroweapons ban is the dual-use nature of the underlying technologies. Unlike chemical weapons precursors, which have limited legitimate applications, neurotechnology has enormous potential for medical treatment, cognitive enhancement, education, and human-machine collaboration. Deep brain stimulation treats Parkinson's disease. EEG headsets enable paralyzed patients to communicate. Transcranial magnetic stimulation offers hope for treatment-resistant depression.

A treaty must distinguish between legitimate medical and civilian applications and prohibited weapons applications. This distinction should be based on consent. Any use of neurotechnology that alters, monitors, or manipulates neural activity without the informed consent of the individual is presumptively prohibited. Medical applications with consent are permitted. Military applications for enhancement of consenting soldiers may be subject to additional regulation. But any application designed to target non-consenting individuals—whether civilians, enemy combatants, or political opponents—is a neuroweapon.

The source material from the Anna Fubini document describes how the same technologies that could "help the emotionally distressed" through brain stimulation are being used to "alter targets moods by stimulating the brain to exhibit certain rhythms (angry, sad, sleepy etc)" without consent. The same satellite systems that could "monitor and stop terrorists" are being used to "spy on political targets" and "terrorise/torture them." The same neurophones that could "communicate with the deaf" are being used to "make the target hear voices."

The Role of the United Nations

The United Nations is the natural forum for negotiating a neuroweapons convention. The UN Human Rights Council has already taken an interest in neurotechnology, passing Resolution 51/3 in 2022 requesting a study on neurotechnology and human rights. The UN Institute for Disarmament Research has already recognized the need for a ban. The 2025 UNESCO Recommendation on the Ethics of Neurotechnology provides an ethical foundation.

The OHCHR call for input, to which the source material responds, demonstrates that the UN system is actively considering the human rights implications of neurotechnology. The source

material recommends that the OHCHR "recognize non-consensual neurotechnological manipulation as a distinct human rights violation" and "develop specialized investigative capacities within human rights mechanisms."

A neuroweapons convention could be negotiated under the auspices of the Conference on Disarmament, or through a special diplomatic conference convened by the UN General Assembly. The precedent of the Ottawa Treaty banning landmines—negotiated outside the traditional disarmament framework—offers an alternative model if the Conference on Disarmament proves unable to act.

The Norm Against Manipulating the Mind

Beyond its legal provisions, a neuroweapons convention would establish a powerful global norm: the human mind is inviolable. No state, no corporation, no individual has the right to access, alter, or manipulate another person's consciousness without their consent. This norm would serve as a foundation for future legal developments, including national legislation, regional frameworks, and customary international law.

The source material from the European Parliament's 1999 resolution captures this principle: the ban is needed because these weapons represent a threat to "human rights and democracy." They attack the very foundation of human autonomy—the ability to think, feel, and decide for oneself. A person whose mind can be remotely manipulated is no longer a free person. A society in which such manipulation is possible is no longer a democracy.

The alternative to a global ban is a future in which neuroweapons become a standard tool of state power, used for covert operations, political repression, and social control. The source material describes this future in stark terms: "The alternative—allowing neurotechnology to develop in a human rights vacuum—risks creating a world where thought itself is no longer private, and the human mind becomes just another battlefield."

The Time to Act

The technologies exist. The victims exist. The legal frameworks exist as models. What is missing is the political will to act. The European Parliament recognized this in 1999: the ban requires "the global pressure of the informed general public on the governments."

That pressure must come from citizens who understand the threat, from legislators who are willing to act, and from the global community that recognizes the stakes. The OHCHR call for input represents an opportunity to make the case for a ban at the highest levels of international human rights governance. The source material urges the OHCHR to "center victim testimony in policy development despite the technical complexity involved" and to "create bridges between human rights frameworks and technical communities to enable effective oversight."

The case for a global ban on neuroweapons is clear. The precedent exists. The urgency is undeniable. The only question is whether the international community will act before the battlefield of the mind becomes a permanent feature of human existence.

CHAPTER 39

An Enforced Moratorium — How Would It Work?

The case for a global ban on neuroweapons is compelling, but a ban without enforcement mechanisms is a paper promise. The Chemical Weapons Convention (CWC) and Biological Weapons Convention (BWC) demonstrate that effective prohibition requires verification, inspection, and consequences for non-compliance. An enforced moratorium on neuroweapons must be built on the same principles, adapted to the unique challenges of technologies that can be miniaturized, concealed, and deployed remotely.

The Verification Challenge

Neuroweapons present a verification problem unlike chemical or biological agents. A chemical weapon leaves residues that can be sampled. A biological agent leaves genetic traces. But a directed-energy neuroweapon—a microwave beam that disrupts neural function, an ultrasonic device that induces cognitive fog, or an electromagnetic system that delivers synthetic telepathy—leaves no physical residue at the point of impact. The weapon itself may be kilometers away, or it may be a device that looks like a cell tower or a streetlamp.

The source material from the TI One Voice legislative reform proposals identifies this challenge directly. The proposed Belgian criminal code amendment addresses it through a "reversal of evidentiary burden regarding the technical feasibility of the alleged device once the victim can demonstrate a pattern of symptoms consistent with published bioeffects literature (referencing standards such as IEEE C95.1 and ICNIRP guidelines for non-thermal effects)." This approach—using symptom patterns and bioeffects literature as the basis for shifting the burden of proof—provides a model for verification at the individual level.

But verification at the international level requires more. The source material proposes "mandatory registration of all directed-energy devices above defined power thresholds with a new EU-level registry managed by the European Defence Agency (EDA) or a designated civilian authority." This registry model, adapted from the CWC's requirement that states declare their chemical weapons stockpiles, would create a baseline for verification. Any device not registered, or any use of a registered device outside declared parameters, would constitute a violation.

The Inspection Mechanism

The CWC operates through a combination of routine inspections and challenge inspections. The source material proposes a similar two-tier approach. At the national level, it calls for "a specialised technical assessment unit within the Comité R, staffed with engineers and physicists with appropriate security clearances, capable of independently auditing electromagnetic emissions from suspected government installations." This unit would conduct routine inspections of declared facilities and respond to complaints.

At the international level, the source material's proposal for an EU regulation would establish "a right to electromagnetic environmental integrity as an extension of the Article 7 Charter right to private life, recognising that deliberate irradiation of a person's home or body constitutes a violation comparable to unlawful surveillance." This right provides the legal basis for challenge inspections—any state that suspects another state of using neuroweapons could request an inspection of the suspected facility.

The inspection teams would need technical capabilities that do not currently exist in most oversight bodies. The source material specifies that these teams must be "equipped with RF spectrum analysis, audiological assessment, and neuroimaging capabilities, to provide independent medical-technical evaluation." These are not standard law enforcement tools. A moratorium on neuroweapons requires investment in the technical infrastructure to detect them.

The Enforcement Pyramid

Enforcement operates at three levels: individual, national, and international. At the individual level, the source material proposes criminal penalties for the use of neuroweapons against persons. The Belgian proposal specifies "penalties ranging from 1 to 10 years imprisonment, with higher ranges for organised conduct and for offences committed by or on behalf of state actors." This creates individual criminal liability, analogous to the CWC's requirement that states criminalize chemical weapons use under their domestic law.

At the national level, the source material proposes "a statutory duty on police zones to log and aggregate complaints of coordinated harassment patterns, even where individual incidents appear minor in isolation." This addresses the problem of fragmented reporting—victims who go to multiple police stations with complaints that are dismissed individually but form a pattern collectively. The aggregation requirement creates a mechanism for identifying systematic abuse.

At the international level, the source material's OHCHR submission proposes "an OHCHR special rapporteur on neurotechnological human rights violations" and "standardized protocols for investigating complaints of technological harassment." These mechanisms would provide international oversight and create a reporting pathway for victims whose national governments are unwilling or unable to act.

The Victim-Centered Approach

The source material emphasizes that enforcement must be victim-centered. The OHCHR submission states: "The administration of justice must evolve to recognize that the absence of visible wounds does not equate to the absence of violation. Neurotechnological torture leaves profound psychological and neurological damage while evading traditional documentation methods."

This requires changes in how evidence is gathered and evaluated. The source material proposes "a specialised diagnostic centre (possibly attached to an existing university hospital) equipped with RF spectrum analysis, audiological assessment, and neuroimaging

capabilities, to provide independent medical-technical evaluation of TI claims." These centers would serve both as evidence-gathering facilities and as support hubs for victims.

The Whistleblower Protection

The source material identifies a critical enforcement gap: the people who know about neuroweapons programs are often bound by secrecy oaths. The OHCHR submission calls for "whistleblower protections for researchers and engineers exposing weaponization." Without these protections, the most credible evidence of neuroweapons programs may never reach enforcement bodies.

The Timeline for Implementation

The source material's legislative proposals provide a realistic timeline. The Belgian parliamentary inquiry commission would produce "a public report with binding legislative recommendations within 18 months." The EU regulation would require "statutory response deadlines (90 days for acknowledgement, 12 months for preliminary findings)" for complaints. These timelines are ambitious but achievable.

The alternative to an enforced moratorium is the status quo: a world in which neuroweapons develop in secrecy, victims suffer without recourse, and the international community lacks the tools to act. The source material's OHCHR submission captures the urgency: "Without immediate intervention, these technologies will continue to enable human rights violations on a global scale while remaining invisible to justice systems."

An enforced moratorium is not a utopian dream. It is a practical necessity, built on existing models, adapted to new challenges, and grounded in the testimony of those who have already experienced the reality of neuroweapons. The question is not whether it can be done, but whether the international community has the will to do it.

CHAPTER 40

How to Detect Illegal Neurotechnology Use

The detection of illegal neurotechnology use presents a unique forensic challenge. Unlike conventional crimes that leave physical evidence at a scene, neurotechnology attacks can occur remotely, leave no visible trace, and employ technologies that official institutions may not acknowledge exist. Yet detection is not impossible. As the CPDE report documents across 1,996 analyzed testimonies from 1,264 global locations, the patterns of neurotechnology abuse are consistent, measurable, and amenable to forensic investigation — provided the right tools and protocols are employed.

Radio Frequency (RF) Spectrum Analysis

The most immediately accessible detection method is radio frequency spectrum analysis. Many neurotechnology systems, particularly those used for remote neural monitoring and directed energy effects, operate by transmitting electromagnetic energy through the air. These transmissions leave spectral signatures that can be captured and analyzed by commercially available spectrum analyzers.

The source material's technical documentation describes "Massive MIMO as a Cognitive Weapon" and other systems that rely on phased array antenna technology to direct energy at specific individuals. These systems operate in licensed and unlicensed frequency bands, often using beamforming techniques that concentrate energy in narrow beams. An RF spectrum analyzer, properly deployed, can detect anomalous signals in the environment — particularly signals that appear to track a person's movement, that exhibit unusual modulation patterns, or that persist in proximity to an individual reporting neurotechnology effects.

The forensic protocol for RF detection requires:

Baseline spectrum recording of the individual's environment when no effects are reported

Continuous monitoring during periods when effects are experienced

Correlation of anomalous signals with reported physical or cognitive effects

Direction-finding to locate the source of detected signals

The UN submission response notes that "no practical shielding or jamming techniques are known to exist at the time of writing," meaning that detection, rather than prevention, is the primary investigative tool available to victims.

Bio-Effect Monitoring

The second pillar of detection is bio-effect monitoring — the systematic measurement of physiological changes that correlate with reported neurotechnology effects. The source material documents a consistent pattern of physical symptoms across thousands of reports:

localized heating, burning sensations, muscle contractions, sleep disruption at precise times, and cognitive interference.

These effects can be measured using:

Thermal imaging cameras to detect localized heating patterns on the skin that correspond to reported directed energy sensations

Electroencephalography (EEG) to monitor for anomalous neural activity patterns during reported cognitive interference

Heart rate variability monitors to track autonomic nervous system responses correlated with harassment episodes

Actigraphy watches to document sleep disruption patterns with precise time-stamping

The CPDE report emphasizes that these measurements must be time-locked — the key forensic signature is not merely that an individual experiences symptoms, but that those symptoms correlate with verifiable external events and measurable physiological changes. The report documents that "Targeted Individuals report that the voice phenomenon, and the broader campaign of intrusion, intensifies and quiets in correlation with verifiable external events: a public statement, a journey, a meeting with a particular person, the publication of certain material online."

Forensic Documentation Protocols

The third essential component is a standardized forensic documentation protocol that captures evidence in a form admissible to courts and human rights bodies. The TI One Voice platform (one.witysk.org), operated by a Brussels-based human-rights initiative, provides a model for this. Its structured digital archive captures:

Geographic precision (specific locations, not general areas)

Temporal data (exact times and dates of incidents)

Incident taxonomy (categorization of the type of effect experienced)

Impact documentation (physical, psychological, social, economic)

Evidence uploads (photographs, recordings, medical reports)

The protocol recommended in the UN submission response includes: "Standardized protocols for investigating complaints of technological harassment" and "forensic methodologies to detect and document neurotechnological attacks." These protocols must be designed to capture evidence that is time-stamped, geolocated, and independently verifiable.

Medical Forensics and the Physical Markers

The source material documents that physical marks and injuries appear in a subset of cases: “Visible marks, injuries, or scarring not resulting from self-infliction or known consensual trauma. Some individuals report indications of foreign objects or marks consistent with non-consensual procedures.”

These physical markers require specialized medical forensic examination. The protocol should include:

Full-body dermatological examination with high-resolution photography

Biopsy of any unexplained lesions or burns

X-ray or MRI to detect foreign objects

Blood work to identify any biomarkers of directed energy exposure

The CPDE report notes that “medical professionals misdiagnose technological torture as mental illness” — a failure that both compounds the victim’s trauma and destroys forensic evidence. Medical forensics must be conducted by professionals trained to recognize the physical effects of neurotechnology, not by clinicians who reflexively reach for psychiatric diagnoses.

The Institutional Infrastructure Gap

The source material is clear that detection is not merely a technical problem but an institutional one. The UN submission response calls for “specialized investigative units within law enforcement” and “inter-agency task forces to consolidate expertise across law enforcement, health systems, and technology ethics.” The CPDE report documents the current reality: “Police refuse investigation… medical professionals reflexively diagnose psychiatric illness… legal systems offer no remedy.”

This institutional failure is itself a human rights violation — the denial of effective remedy under ICCPR Article 2(3). The source material documents specific locations where this failure has occurred: the 4th Police District in the United States, Lokale Politie Zone Zuid in Belgium, various precincts in London and Paris. These are not abstract failures but concrete, documented instances of institutional refusal to engage with the evidence.

The Path to Verification

Detection of illegal neurotechnology use is possible. The tools exist: RF spectrum analyzers, thermal imaging cameras, EEG monitors, actigraphy watches, and standardized forensic documentation protocols. The evidence patterns are documented across thousands of cases from hundreds of locations worldwide. What is missing is not the technology of detection but the institutional will to deploy it.

The CPDE report’s conclusion applies directly to the detection challenge: “The scale, consistency, and severity demand coordinated multi-sectoral response.” Detection requires law enforcement agencies equipped with RF analysis capabilities, medical systems trained to

recognize technological effects, legal frameworks that accept this category of evidence, and international bodies willing to investigate complaints that cross borders.

Without this institutional infrastructure, detection remains the province of individual victims armed with consumer electronics — a situation that the source material documents as producing systematic failure. With it, the consistent patterns documented across 1,264 locations can become the basis for investigation, accountability, and ultimately, protection of the last private place: the human mind.

CHAPTER 41

Proving Mind Control After the Fact — Forensic Challenges

The evidentiary challenge at the heart of neuro-rights enforcement is this: how does a court determine that a person was under the influence of mind control? The question is not hypothetical. Thousands of individuals across 1,264 global locations have submitted structured testimonies describing experiences they attribute to non-consensual neurotechnology use. The CPDE report documents the consistency of these accounts. Yet no prosecution for neurotechnology-facilitated torture, assault, or privacy violation has succeeded anywhere in the world. The forensic gap is the reason.

The Nature of the Evidentiary Problem

Mind control, as described in the source material, does not leave the kind of evidence that conventional forensic science is designed to detect. There are no syringe marks, no blunt-force injuries that match a weapon profile, no DNA samples from a perpetrator who was physically present. The alleged mechanism is remote, wireless, and often intermittent. The effects are neurological and psychological rather than structural. The victim may appear physically unharmed while reporting experiences that, if genuine, constitute severe suffering.

The CPDE report identifies the core difficulty: “The absence of visible wounds does not equate to the absence of violation.” But courts operate on visible evidence. A judge cannot see a directed-energy beam. A jury cannot examine a neural signal transmitted through walls. The technology, if it exists as described, operates below the threshold of ordinary perception and leaves no physical trace that standard forensic protocols are designed to recover.

This creates a perverse evidentiary situation. The more sophisticated the alleged attack — the more purely neurological its mechanism — the less conventional evidence it leaves behind. A beating leaves bruises. A stabbing leaves a wound. A directed-energy neuroweapon, if it functions as reported, leaves only the victim’s testimony and whatever physiological markers a trained examiner knows to look for.

Biomarkers: The Search for Physical Evidence

The source material documents specific physical effects that Targeted Individuals report with consistency: skin lesions and burns appearing without explanation, localized heating sensations, forced muscle contractions, cardiac arrhythmias, sleep disruption patterns that track external operational factors rather than internal pathology. These effects, if systematically documented and analyzed, could constitute biomarkers of neurotechnology exposure.

The challenge is that no biomarker is currently validated for this purpose. The CPDE report notes that medical examinations of physical marks have been conducted in individual cases,

but no standardized protocol exists for distinguishing a directed-energy injury from a dermatological condition of unknown origin. The same skin lesion could be a burn from a covert device or an idiopathic rash. The same cardiac symptom could be induced by electromagnetic means or by anxiety. Without a validated biomarker — a measurable physiological change that is reliably produced by the alleged technology and not produced by other causes — the physical evidence remains ambiguous.

The path forward is epidemiological. The CPDE dataset of nearly 2,000 structured testimonies provides statistical power to identify patterns that would not be visible in individual cases. If a specific constellation of symptoms — localized heating at particular body sites, sleep disruption at particular circadian phases, particular types of skin lesions — appears across thousands of independent reports from hundreds of locations, that pattern constitutes evidence. It is not proof of causation, but it is evidence that something systematic is occurring, and it provides a basis for targeted forensic investigation.

Witness Testimony: The Challenge of Credibility

The most direct evidence in any mind-control case is the testimony of the alleged victim. But that testimony faces a credibility barrier that is nearly insurmountable under current legal frameworks. The source material documents the reflexive diagnostic response: when a person reports hearing voices or experiencing persecution, the default institutional reaction is to classify the experience as psychosis.

The TI One Voice Declaration's analysis of eight lines of evidence demonstrates why this classification is scientifically untenable in the cases documented. The late, sudden onset in previously well-functioning adults. The voices with stable, distinguishable personalities that possess verifiable external information the victim could not have known. The banal surveillance content rather than grandiose delusion. The quieting periods uncorrelated with antipsychotic medication. The time-locked correlations with verifiable external events.

But these arguments are complex. They require a court to understand the difference between primary psychosis and technologically induced experience — a distinction that most clinicians do not currently recognize. The victim who testifies to being subjected to voice-to-skull harassment will be heard by judges, jurors, and lawyers who have been trained to associate exactly such reports with mental illness. The credibility barrier is not a failure of the evidence. It is a failure of the institutional framework to distinguish between two phenomenologically different categories of experience.

Pattern Analysis: The Aggregate Case

The most powerful forensic tool for proving mind control after the fact may not be individual biomarkers or individual testimony, but pattern analysis across cases. The CPDE report documents that reports from Brussels and Auckland, from Paris and Pyongyang, from Miami and Mariupol, converge on remarkably specific features: the same kinds of voice phenomena, the same operator behaviors, the same architectural descriptions of the alleged apparatus, the same sequencing of phases over time.

This convergence is structural, not stylistic. It is not that thousands of people independently describe feeling watched. It is that thousands of people independently describe the same specific sequence of technological intrusion: initial surveillance, then voice phenomena with particular characteristics, then physical effects at particular body sites, then social harassment by coordinated strangers, then institutional dismissal when they seek help. The specificity of the pattern, documented across cultures and languages, is itself evidence.

The CPDE report states this directly: “To dismiss this convergence as the coincidental product of thousands of independent identical hallucinations is to require a coincidence whose probability, on any reasonable account, is much lower than the probability of the convergence reflecting a real external common cause.” A court that hears expert testimony on the statistical significance of this pattern has a basis for evaluating individual claims that does not depend solely on the credibility of any single witness.

The Legal Framework Gap

No current legal framework contains provisions for proving non-consensual neurotechnology use. The UN report on neurotechnology and human rights (A/HRC/57/61) identified this gap but did not propose specific forensic protocols. The Chilean constitutional reform protects brain rights but does not specify how violations would be proven. The Colorado and California neurodata laws focus on data protection, not on proving that neural manipulation occurred.

The source material’s call for “forensic methodologies to detect and document neurotechnological attacks” and “standardized protocols for investigating complaints of technological harassment” identifies the specific legislative and institutional reforms needed. Until those protocols exist, the evidentiary challenge remains insurmountable. Courts will continue to dismiss claims for which no recognized forensic methodology exists, and perpetrators will continue to operate with impunity.

The Path Forward

Proving mind control after the fact requires a multi-layered forensic approach. Biomarker validation through controlled studies comparing documented cases with matched controls. Standardized documentation protocols that capture the specific physiological and phenomenological patterns identified in the CPDE dataset. Expert testimony on the statistical significance of cross-cultural convergence. Legal recognition that the absence of conventional physical evidence does not equate to the absence of harm.

The source material documents that the tools for this forensic approach already exist: RF spectrum analyzers, thermal imaging, EEG monitoring, actigraphy watches, structured incident logs. What is missing is the institutional framework that validates these tools for forensic purposes and the legal framework that accepts this category of evidence.

The CPDE report’s conclusion applies directly to the forensic challenge: the scale, consistency, and severity of the documented phenomenon demand a coordinated multi-sectoral response. That response must include the development of validated forensic

protocols for detecting and proving non-consensual neurotechnology use. Without them, the right to mental privacy remains a paper right — enforceable in theory, unenforceable in practice.

CHAPTER 42

Defining a Victim of Neurotechnology — Recognition and Support

The Invisible Victim

A crime occurs. The victim experiences physical pain, psychological torment, and the systematic destruction of their private life. Yet when they seek help — from police, from doctors, from lawyers — they are told nothing happened. The absence of visible wounds, the lack of a weapon that fits legal definitions, the very strangeness of the reported mechanism, all conspire to erase the victim's experience. They become not a victim but a problem: a person with a story too difficult to believe, too uncomfortable to investigate, too threatening to the established order.

This is the reality of neurotechnology abuse. The victims exist. Their suffering is documented across thousands of testimonies, in medical records, in police reports that went nowhere, in the quiet desperation of lives destroyed by forces that leave no trace visible to conventional forensics. But without a legal framework that recognizes their victimhood, they remain invisible — denied the most basic right of any person who has been harmed: the right to be believed.

Defining a Victim of Neurotechnology

The TI One Voice Declaration v3.1b provides the first formal definition. A Targeted Individual is a "natural person who experiences ongoing, systematic, organized harassment and stalking through advanced technologies, often conducted covertly and without their consent." The definition is careful and precise. It does not require the victim to prove the mechanism of harm. It requires only that they report experiencing the characteristic pattern of intrusion: remote and wireless attacks that violate privacy, physical integrity, and psychological well-being, with no known shielding or jamming techniques available.

The core manifestations, as documented in the Declaration and corroborated across the CPDE dataset of nearly 2,000 structured testimonies, include:

Voice-to-Skull (V2K): The perception of voices, messages, or auditory phenomena that appear to originate externally and are not under the individual's control. These voices are described as having stable, distinguishable personalities — not the flickering, disorganized voices of primary psychosis, but consistent identities that retain verbal mannerisms, opinions, and behavior over months and years.

Directed Energy Effects: Physical sensations of pain, burning, vibration, pressure, or other somatic effects that individuals assert are induced externally. The CPDE dataset documents skin lesions, burns, localized heating, forced muscle contractions, and systematic sleep deprivation.

Physical Marks and Injuries: Visible marks, scarring, or indications of foreign objects not resulting from self-infliction or known consensual trauma.

Psychological and Behavioral Effects: Symptoms of induced fear, anxiety, sleep deprivation, or attempted behavioral influence.

Systematic Interference: Ongoing disruption of privacy, sleep, or normal functioning through intrusive means asserted to be external and covert.

The definition is deliberately cause-neutral. It does not require the victim to prove the technological mechanism. It does not require the victim to identify the perpetrator. It requires only that the pattern of reported experiences meets the established criteria. This is consistent with how other victim categories are defined in international law: the victim of torture does not need to name the torturer to be recognized as having suffered torture.

Why This Is Not Mental Illness

The most common response to a Targeted Individual's report is the reflexive diagnosis of schizophrenia or delusional disorder. The reflex is understandable — mainstream psychology has trained clinicians to reach for these categories when someone reports hearing voices or experiencing persecution. But the categories do not fit. The CPDE report documents eight independent lines of evidence that demonstrate why the standard psychiatric dismissal is untenable.

First, onset is late and sudden in previously well-functioning adults. Schizophrenia typically emerges between adolescence and the late twenties. Targeted Individuals are often in their thirties, forties, or fifties with no prior psychiatric history and no significant prodromal warning signs. Second, the voices reported have stable, distinguishable personalities — not the flickering symptoms of disordered processing, but consistent identities that can be characterized in detail. Third, and most decisively, the voices possess verifiable external information — specific facts the target had no way of knowing in advance, subsequently confirmed by independent means. A hallucination cannot, in principle, contain information not already in the head of the person hallucinating.

Fourth, the reported content is banal rather than grandiose. Primary psychotic delusions tend toward the grandiose or mythologically framed. Targeted Individuals report voices that comment on daily routines, reference the recent grocery shop, mock ordinary habits — the phenomenology of continuous low-grade surveillance, not a brain in chemical disarray. Fifth, the target's ability to characterize the apparatus in detail — producing careful, multi-year, internally consistent accounts — demands the preserved cognitive function that the diagnostic categories are supposed to be precluding.

Sixth, convergent reports across thousands of independent witnesses worldwide produce structural, not stylistic, convergence on remarkably specific features. Seventh, quieting periods are uncorrelated with antipsychotic medication but track external operational factors. Eighth, time-locked correlations with verifiable external events produce symptom intensifications that a primary psychotic illness could not produce.

Taken together, these eight observations make the standard psychiatric diagnosis untenable. The working hypothesis that best accounts for all eight — that some combination of surveillance technologies, coordinated human operators, and cumulative harassment

produces the reported experience — requires none of the observations to be set aside as inexplicable.

The Human Rights Framework

The experiences documented by Targeted Individuals constitute systematic violations of fundamental human rights protected under international law. The right to privacy, enshrined in Article 12 of the Universal Declaration of Human Rights and Article 17 of the International Covenant on Civil and Political Rights (ICCPR), prohibits arbitrary interference with privacy, family, home, and correspondence. The CPDE dataset reveals that the most frequently cited locations of intrusion are not public spaces but the most intimate zones of life: bed, kitchen, apartment, domicile, house. The sanctity of the home — the foundational privacy zone recognized in every legal system — is reported as completely violated.

ICCPR Article 7 and the Convention Against Torture establish an absolute prohibition against torture and cruel, inhuman, or degrading treatment. The reports from Targeted Individuals meet or approach each of the four legal criteria: severe physical or mental suffering, intentionally inflicted, for a purpose such as intimidation or punishment, and with official involvement or acquiescence. The physical effects — unexplained burns, localized heating, forced muscle contractions, systematic sleep deprivation — constitute severe suffering if genuine. The psychological dimension — abusive, threatening, controlling voices perceived as externally transmitted — creates a state of psychological torture.

The right to the highest attainable standard of health (International Covenant on Economic, Social and Cultural Rights Article 12) is violated by the systematic destruction of physical and mental health. The right to an effective remedy (ICCPR Article 2(3)) is denied by the universal pattern of institutional dismissal documented across the corpus.

The Need for Official Recognition

The most urgent need is official recognition of the victim category. Without it, every other protection mechanism remains inaccessible. Recognition means:

Legal standing: The victim can file complaints, seek restraining orders, and pursue civil remedies through mechanisms that accept their category of harm as actionable.

Access to investigation: Law enforcement agencies have protocols for receiving and investigating reports of neurotechnology abuse, rather than immediately referring the reporter to mental health services.

Medical recognition: Healthcare systems have diagnostic codes and treatment protocols for technology-facilitated harm, distinguishing it from primary psychiatric illness.

Social support: Victim support services — crisis lines, trauma therapy, practical advocacy — are available without requiring the victim to prove the mechanism of harm.

Data collection: Official statistics track the prevalence and distribution of reported neurotechnology abuse, enabling public health responses and policy development.

The TI One Voice Declaration provides the framework for this recognition. Its Article II characterizes the violations as criminal acts under existing law — assault, battery, stalking, harassment, torture — and as human rights violations under international law. Its Article III demands victim recognition, government support and protection, legislative action, and public disclosure. The Declaration is both a legal document and a political instrument: a petition for investigation, recognition, protection, and remediation.

The Path Forward

The CPDE report documents that the trajectory of a Targeted Individual is predictable and devastating: initial confusion, followed by recognition and fear, then seeking help only to encounter dismissal, leading to social and economic collapse, and finally isolation and crisis. Each stage represents a missed opportunity for intervention. Breaking this cascade requires action at every point:

Immediate: Formal acknowledgment that a substantial population reports consistent, severe experiences that current systems fail to address. Specialized reporting protocols within law enforcement. Cause-neutral victim support services.

Medium-term: Legislative change expanding anti-stalking and harassment laws to address coordinated harassment by networks. Inter-agency task forces consolidating expertise. Funded academic research.

Long-term: International treaty frameworks governing civilian use of directed-energy devices and neural technologies. Technical standards body governance for dual-use devices. Civil society monitoring and advocacy capacity.

The hardest question — what causes this shared reality — requires the coordinated inquiry that recognition is designed to catalyze. What is no longer acceptable is silence, dismissal, and the systematic denial of human rights to thousands of people across every inhabited continent. The victims exist. Their suffering is documented. The only question is whether the world will choose to see them.

PART VII

The Ethics of Neurotechnology

The question at the heart of neurotechnology ethics is deceptively simple: when, if ever, is it acceptable to access, monitor, or alter another person's brain activity without their consent? The answer requires navigating a terrain where absolute prohibitions meet genuine emergencies, where individual autonomy confronts collective necessity, and where the boundaries of personhood itself are at stake.

The Absolute Baseline

International human rights law establishes clear red lines. The prohibition of torture and cruel, inhuman or degrading treatment — as articulated in the UN report A/HRC/57/61 — is absolute. Neurotechnologies used to cause suffering, to extract confessions, or to degrade human dignity can never be justified, under any circumstances, even during states of emergency. The report explicitly identifies "neurocorrectives" as inherently degrading treatment, and warns that brain-reading techniques applied to detained persons constitute prohibited pressure under the right against self-incrimination (A/HRC/57/61, paras. 34–36). These are not matters for ethical balancing; they are categorical prohibitions.

The Medical Exception

The most defensible case for non-consensual neurotechnology use arises in medicine. Emergency treatment for epilepsy, locked-in syndrome, or acute psychiatric crisis may require intervention when the patient cannot consent. The UN report acknowledges that neurotechnologies hold "great potential to improve people's mobility and autonomy" in medical contexts (A/HRC/57/61, para. 1). Yet even here, the ethical terrain is treacherous.

Deep brain stimulation patients have reported "a changed sense of agency and identity" (A/HRC/57/61, para. 27). The line between care and coercion blurs when treatment alters personality, memory, or self-perception.

The critical safeguard is proportionality. Non-consensual intervention must be time-limited, minimally invasive, and subject to independent oversight. The right to refuse treatment — to choose one's own mental state — remains fundamental. States must guarantee "the right to choose or refuse the treatment they want with the full knowledge of the risks and benefits" (A/HRC/57/61, para. 28).

The Criminal Justice Dilemma

Proposals to use neurotechnology for evidence-gathering in criminal justice present profound ethical challenges. The UN report is unequivocal: brain-reading techniques in investigations violate the presumption of innocence and the right not to testify against oneself (A/HRC/57/61, para. 36). Subjecting accused persons to neural surveillance — even with court authorization — constitutes "a form of prohibited pressure."

Yet the ethical calculus shifts when neurotechnology might prevent imminent harm. Could a brain-computer interface detect intent to commit violence? The report warns that neurotechnologies risk "the criminalization of mere thought" — punishing people for unexpressed intentions before any action occurs (A/HRC/57/61, para. 20). This violates the foundational principle that thought, however dangerous, is not crime.

The Slippery Slope

The greatest ethical danger may not be deliberate abuse but incremental normalization. Consumer neurotechnology companies already collect unknown quantities of neural data, stored for future decoding (A/HRC/57/61, para. 25). Workplace monitoring, educational surveillance, and neuromarketing each represent a step along a slope that descends from voluntary use to coercive expectation. The UN report documents how neurogaming devices can drive "compulsive or addictive use" by targeting neural reward circuits (A/HRC/57/61, para. 21).

Once non-consensual use is accepted in any context, the ethical framework weakens. The burden of proof must remain on those who would override consent. And that burden is extraordinarily heavy — because the inner realm of thought, the *forum internum*, is the last sanctuary of the autonomous self.

CHAPTER 43

The Absolute Prohibition — Torture and Neurotechnology

The prohibition of torture is the most absolute prohibition in international human rights law. No exceptions, no emergencies, no national security justifications. And now, the United Nations has made clear that non-consensual neurotechnology use that causes suffering falls squarely within this prohibition.

The UN Human Rights Council Advisory Committee's 2024 report (A/HRC/57/61) is unambiguous: "The absolute character of the prohibition of torture, cruel, inhuman or degrading treatment or punishment underlines the gravity of the interferences" posed by neurotechnologies (para. 34). This is not a matter of debate or interpretation. When a neurotechnology device is used to cause severe mental or physical suffering, it constitutes torture under international law.

Neurotechnology as Torture

How can a device become an instrument of torture? The UN report identifies several mechanisms. First, "brain-reading" techniques used against detained persons constitute "an inherently degrading treatment in all circumstances" (para. 34). Forcing someone to undergo neural surveillance — to have their thoughts extracted, their memories scanned, their intentions decoded — violates the core of human dignity.

Second, "brain-writing" techniques — neuromodulation that alters mental states without consent — can inflict profound psychological suffering. The report specifically references "neurocorrectives," which it describes as "an inherently degrading treatment in all circumstances" (para. 34). These are devices or procedures applied to modify behavior or cognition against a person's will, often justified as "treatment" or "correction."

Third, medical or scientific experimentation without free consent is absolutely prohibited. The report states: "Medical, scientific or biological experimentation without the free consent of the person concerned is prohibited under international human rights law" (para. 34). Special protection is required for persons not capable of giving valid consent, particularly those deprived of their liberty.

The Absolute Nature of the Prohibition

The prohibition of torture is non-derogable. It applies even during states of emergency, armed conflict, or any other public crisis. The UN report emphasizes that "no statements or confessions or other evidence obtained in violation of article 7 of the International Covenant on Civil and Political Rights may be invoked in any proceedings, not even during a state of emergency" (para. 35).

This means that evidence obtained through coercive neurotechnology — whether brain-reading to extract confessions or neuromodulation to induce compliance — is inadmissible in

any legal proceeding. The burden falls on the State to prove that any statements made by accused persons were given "of their own free will" (para. 35).

The Vulnerability of Detained Persons

The UN report specifically warns about the vulnerability of persons deprived of liberty. "Subjecting detained or imprisoned persons to 'brain-reading' techniques is clearly a form of prohibited pressure" (para. 36). This applies regardless of the alleged crime, the perceived threat, or the urgency of the investigation.

The report also notes that "persons with disabilities are particularly exposed to medical and scientific procedures, research and experimentation" (para. 38). Neurotechnologies developed without proper human rights guidelines expose them to "a disproportionate risk of physical and/or mental harms" (para. 38). Without meaningful engagement and participation, these technologies risk "contributing to stigmatization and objectification, reinforcing structural-based discrimination" (para. 38).

The Convention on the Rights of Persons with Disabilities

The UN report explicitly references the Convention on the Rights of Persons with Disabilities as providing "a protective framework that cannot be departed from" (para. 39). Article 15 of the Convention prohibits torture and cruel, inhuman or degrading treatment or punishment, specifically including medical or scientific experimentation without free consent.

This is particularly relevant as neurotechnology companies increasingly target persons with disabilities for clinical trials and device implantation. The report warns that "consent to participate in trials may be provided under coercion or without being fully informed or made aware of negative side-effects on, for example, health, privacy and autonomy" (para. 37).

The Duty to Prevent

States have an affirmative duty to prevent torture through neurotechnology. This requires not only prohibiting direct use by state actors but also regulating private companies that develop and deploy these technologies. The UN report calls for "explicit prohibitions of unsolicited exposure to neuromodulation procedures" (para. 29).

The report also emphasizes that "domestic law must ensure that statements or confessions obtained in violation of this prohibition are excluded from the evidence" (para. 35). This is not optional — it is a binding obligation under international law.

No Justification, No Exception

The prohibition of torture admits no justification. Not national security. Not medical necessity. Not the prevention of crime. Not the demands of military operations. The UN report is clear: "coercive uses of neurotechnologies to augment resilience on the battlefield (i.e. to reduce emotions, empathy or the need for sleep) should be prohibited" (para. 29).

This absolute prohibition extends to all forms of non-consensual neurotechnology use that cause severe mental or physical suffering. The inner realm of thought — the *forum internum* — is inviolable. And those who would breach it through coercion, deception, or force commit an act that the international community has deemed the gravest of human rights violations.

CHAPTER 44

Medical Exceptions — When Is Non-Consensual Use Ethically Justified?

The absolute prohibition against non-consensual neurotechnology use that causes suffering is clear. But what about cases where the intent is not to harm, but to heal? This chapter examines the most ethically complex scenarios: emergency treatment, interventions for patients who cannot consent, and the thin line between care and coercion.

The Principle of Informed Consent

The UN report (A/HRC/57/61) reaffirms that "ensuring that decisions in neurotechnology interventions and clinical trials are genuinely autonomous is vital" (para. 28). In the medical sphere, human rights standards require that States guarantee people's "right to choose or refuse the treatment they want with the full knowledge of the risks and benefits" (para. 28, citing Committee on Economic, Social and Cultural Rights, general comment No. 25).

This foundational principle applies to all neurotechnology interventions. The report emphasizes that "human rights standards and limitations in relation to medical practices and coercive health treatments, particularly in psychiatry, also apply to neurointerventions" (para. 28).

Emergency Treatment: The Locked-In Syndrome Dilemma

Consider a patient with locked-in syndrome — fully conscious but unable to move or communicate. A brain-computer interface (BCI) could restore communication, but the patient cannot consent. Is it ethical to implant the device without consent?

The UN report acknowledges that neurotechnologies "hold great prospects in protecting mental health and granting personal autonomy" (para. 31). However, it also warns that "persons in vulnerable situations are particularly exposed to medical and scientific procedures, research and experimentation" (para. 38). The Convention on the Rights of Persons with Disabilities provides "a protective framework that cannot be departed from" (para. 39).

The ethical framework requires balancing the potential benefit against the risk of harm. The report notes that "many of those technologies are still experimental" and "the potential long-term effects of using neurotechnologies are largely unknown" (para. 31). For locked-in patients, the benefit — restoring communication — may outweigh risks, but only if the intervention is genuinely life-saving and no less restrictive alternative exists.

Dementia and Cognitive Decline

For patients with severe dementia who cannot provide consent, neurotechnology might slow cognitive decline or manage behavioral symptoms. The UN report notes that "negative

impacts on mental health of children, older persons and persons with disabilities should be considered together with the opportunities" (para. 33).

The critical question is whether the intervention serves the patient's best interests or the convenience of caregivers. The report warns that "enhanced mechanisms should be adopted to ensure that persons in vulnerable situations are effectively protected from all forms of coercion in the context of medical research and applications, and beyond" (para. 29).

Epilepsy and Closed-Loop Systems

Closed-loop neurostimulation systems can detect and prevent epileptic seizures before they occur. These devices monitor brain activity and deliver electrical stimulation automatically, without the patient's conscious awareness. For patients with intractable epilepsy who cannot consent, is implantation justified?

The UN report distinguishes between "reading" and "writing" neurotechnologies (para. 6). Closed-loop systems do both — they read neural signals and write corrective stimulation. The report notes that such systems may "bypass the individual's conscious control or awareness" (para. 5(c)). For epilepsy, the potential benefit — preventing life-threatening seizures — may justify intervention, but only with rigorous safeguards.

The Thin Line Between Care and Coercion

The UN report identifies a fundamental tension: "Neurotechnologies and their applications should first be examined against the foundational principle of human dignity" (para. 17). This requires assessing whether applications "contribute to the objectification and instrumentalization of individuals" and "whether autonomy may be neglected by altering the foundation of their mental processes and states" (para. 17).

The report warns that "a human rights-based approach also requires limiting the implementation of applications that are unable to comply with it" (para. 17). This means that even well-intentioned medical interventions must be scrutinized for their potential to undermine human dignity.

The Slippery Slope

What begins as medical necessity can slide into coercion. The UN report notes that "consent to participate in trials may be provided under coercion or without being fully informed or made aware of negative side-effects on, for example, health, privacy and autonomy" (para. 37). This risk is particularly acute for "groups experiencing structural discrimination and marginalized people or those in positions of socioeconomic disadvantage" (para. 37).

The report also warns that "developers may exploit their economic, physical or psychological vulnerabilities or those related to their work situation or age, among others" (para. 37). This applies not only to clinical trials but to any context where power asymmetries exist.

The Need for Safeguards

The UN report calls for “enhanced mechanisms” to protect vulnerable persons (para. 29). These should include:

Independent ethical review for all non-consensual neurotechnology interventions

Clear criteria for what constitutes a genuine medical emergency

Mechanisms for ongoing consent as the patient’s capacity changes

Prohibition of interventions that serve institutional convenience rather than patient welfare

Explicit prohibitions of “unsolicited exposure to neuromodulation procedures” (para. 29)

Conclusion: No Easy Answers

The medical exception to non-consensual neurotechnology use is not a blank check. The UN report is clear: even in medical contexts, “respect for that overarching principle [of human dignity] should be monitored throughout the life cycle of neurotechnologies” (para. 17). The burden of proof falls on those who would intervene without consent, and the threshold for justification is high.

As the report states, “a human rights-based approach also requires limiting the implementation of applications that are unable to comply with it” (para. 17). This means that some neurotechnology applications — however promising — may simply be incompatible with fundamental human rights, even in medical contexts. The line between care and coercion is thin, and crossing it requires not just medical necessity, but rigorous ethical justification, independent oversight, and unwavering commitment to human dignity.

CHAPTER 45

Criminal Justice — Brain-Reading as Evidence

The prospect of using neurotechnology to read minds in criminal proceedings represents one of the most profound challenges to fundamental principles of justice. The United Nations Human Rights Council Advisory Committee, in its landmark 2024 report (A/HRC/57/61), directly addressed this issue, finding that neurotechnologies "being proposed for use in investigation and fact-finding in the criminal domain are highly problematic" (para. 36). The report's conclusion is stark: it is "difficult to see how they can be used without violating the right to due process of law, which applies at all times" (para. 36).

The Right Against Self-Incrimination

At the heart of the matter lies the privilege against self-incrimination — the right not to be compelled to testify against oneself. The UN report is unequivocal: "The presumption of innocence and the right not to testify against oneself require that no direct or indirect physical or undue psychological pressure be exerted by the authorities on the accused to confess or testify" (para. 36).

Brain-reading technologies, whether through fMRI-based lie detection, EEG-based memory detection, or other neural decoding methods, fundamentally bypass the conscious decision to speak or remain silent. They extract information directly from neural activity, without the individual's active participation or control. The UN report identifies this as a form of "prohibited pressure," stating that "subjecting detained or imprisoned persons to 'brain-reading' techniques is clearly a form of prohibited pressure" (para. 36).

This conclusion flows from the absolute nature of the prohibition on coercion in criminal proceedings. The report notes that "no statements or confessions or other evidence obtained in violation of article 7 of the International Covenant on Civil and Political Rights may be invoked in any proceedings, not even during a state of emergency" (para. 35). Domestic law, the report insists, "must ensure that statements or confessions obtained in violation of this prohibition are excluded from the evidence" (para. 35).

The Criminalization of Mere Thought

Beyond the self-incrimination problem, the UN report warns that neurotechnology in criminal justice could lead to "the criminalization of mere thought" (para. 20). The report explains that "with the support of neurotechnologies, people can be punished for unexpressed thoughts or intentions and pre-emptive decisions affecting their lives or interest can be taken even before they translate into action" (para. 20).

This prospect directly threatens the presumption of innocence, which requires that guilt be proven through voluntary actions, not through inferred mental states. The report specifically warns that "violations can happen in particular in the context of counter-terrorism policies if neurotechnologies are used to pre-emptively punish non-violent action" (para. 20). The

ability to detect neural signatures associated with intent or planning could be used to justify preventive detention or enhanced surveillance based on thoughts alone — a concept fundamentally incompatible with liberal democratic legal systems.

The Reliability Problem

Even setting aside the rights-based objections, the scientific reliability of brain-reading for evidentiary purposes remains unproven. The UN report acknowledges that "most neurotechnologies are still being tested and do not provide accurate results in the decoding of people's thoughts, which means that 'mind-reading' or 'mind-steering' are not still achievable" (para. 7). Neuroimaging can "be employed to make inferences about individuals on a large spectrum of mental states, including memories, semantic knowledge, emotions, dreams, inner speech and intentions" (para. 7), but these inferences are probabilistic, not deterministic.

The risk of false positives — where neural activity is misinterpreted as indicating guilt, deception, or criminal intent — is substantial. The report warns that "potential misuses of such information are a major concern" (para. 7), particularly when combined with artificial intelligence, which may introduce "algorithmic bias and lack of traceability and transparency" (para. 7).

A Clear Prohibition

The UN report's framework leaves little room for compromise. While acknowledging that neurotechnologies "are highly problematic" in the criminal domain, the report does not suggest a complete ban on all uses. Rather, it establishes a clear principle: any use of brain-reading in criminal proceedings must not violate the right against self-incrimination, the presumption of innocence, or the prohibition on torture and cruel treatment.

The burden falls on the state to prove that any evidence obtained through neurotechnology was "given of their own free will" (para. 35). Given the nature of these technologies — which can bypass conscious control and extract information without the individual's active participation — meeting this burden will be extraordinarily difficult, if not impossible, in most contexts.

The UN report's conclusion is a warning to legislatures and courts: the use of brain-reading as evidence in criminal proceedings is not a technical question of evidentiary admissibility, but a fundamental challenge to the architecture of procedural justice. The right to remain silent, the presumption of innocence, and the prohibition on compelled self-incrimination are not outdated concepts — they are the bedrock of a just legal system, and they must be preserved against the allure of technological shortcuts.

CHAPTER 46

The Slippery Slope — From Consent to Coercion

The most insidious threat posed by neurotechnology is not the overtly malicious device, but the gradual, seemingly reasonable expansion of its use from voluntary to coerced applications. The UN report (A/HRC/57/61) identifies this dynamic explicitly, warning that the "increased commercialization for non-medical purposes raises important challenges" (para. 2) precisely because the boundary between consent and coercion is porous and easily eroded.

The Workplace: From Wellness to Surveillance

Consider the non-invasive EEG headset marketed to improve employee focus and well-being. A worker voluntarily agrees to wear it during a trial period, receiving feedback on their concentration levels. This appears benign—even beneficial. But the report notes that neurotechnology devices "may be used, even on a permanent basis, while performing different activities" (para. 10). Once the device is embedded in workplace culture, the voluntary choice becomes a de facto requirement. An employee who refuses may be seen as uncommitted, less productive, or hiding something.

The UN report warns that neurodata can be used "to discriminate among people on the basis of their mental signatures, leading to decisions (i.e. in relation to the workplace or insurance coverage) that may have a negative impact on the interests and rights of the data subject" (para. 25). What begins as an optional productivity tool transforms into a mechanism for performance evaluation, promotion decisions, and even termination. The worker who initially consented to "wellness monitoring" finds themselves subject to continuous surveillance of their cognitive states—fatigue, distraction, emotional responses—without the ability to withdraw without penalty.

Educational Surveillance: The Normalization of Monitoring

A similar trajectory unfolds in education. Schools adopt neurotechnology devices marketed as tools for "personalized learning" or "attention monitoring." The UN report notes that "non-invasive neurotechnologies are used for a variety of medical purposes, but the same devices are also being commercialized for different applications in the area of well-being, in the workplace or as an educational tool" (para. 10). A student wears a headband that tracks their focus during class, ostensibly to help them improve concentration.

But once such monitoring becomes routine, the line between support and surveillance blurs. Students who show "inappropriate" neural patterns—daydreaming, emotional distress, or critical thoughts about the institution—may face consequences. The report warns that neurotechnologies "can be used to interfere and manipulate individuals" (para. 21), and that "neurogaming devices can also drive compulsive or addictive use, particularly when neural reward circuits target altering consumption behaviours" (para. 21). The classroom becomes a

laboratory for behavioral modification, where the student's inner mental states are no longer private but subject to institutional evaluation.

Social Scoring: The Ultimate Slippery Slope

The most extreme manifestation of this slide from consent to coercion is the emergence of neural social scoring systems. The UN report raises the alarm that "with the wide deployment of brain-computer interfaces, requiring the collection, processing and storage of massive quantities of personal information, the risk of misuse of data, repurposing, predatory practices and hacking will become significant" (para. 24). Once neural data is collected—whether from workplace devices, educational tools, or consumer headsets—it can be aggregated, analyzed, and used to create "neuroprofiles" that score individuals on cognitive performance, emotional stability, or even political inclinations.

The report explicitly warns that "neuroprofiles may be used to discriminate among people on the basis of their mental signatures" (para. 25). This is not science fiction. The data infrastructure for such scoring is already being built by consumer neurotechnology companies that "are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future" (para. 25). A citizen's neural profile—initially created with their consent for a specific purpose—can be sold to insurers, employers, or government agencies without their knowledge.

The Erosion of Informed Consent

The UN report identifies a fundamental problem: "it is hard to see how users can provide informed consent when their potential long-term effects are largely unknown" (para. 30). Neurotechnology users "may not fully understand the implications of using neurotechnologies" and "may unknowingly surrender control over their cognitive processes" (para. 30). The technology itself compounds this problem because "neurotechnologies allow the level of conscious reasoning to be bypassed, leaving individuals without psychological protection from having their mind involuntarily read" (para. 30).

This creates a paradox: the very act of consenting to neurotechnology use may be compromised by the technology's ability to bypass conscious control. A person who agrees to wear a neural monitoring device may not realize that the device can detect thoughts, emotions, or intentions that they have not consciously expressed. The consent they gave was for one purpose, but the technology's capabilities extend far beyond what they understood.

The Role of Power Asymmetries

The report emphasizes that "in contexts characterized by power asymmetries, for example in the military domain, human rights concerns may arise with regard to interventions for enhancement purposes" (para. 29). The same logic applies in any hierarchical setting—employment, education, healthcare, or criminal justice. When one party holds power over another, the line between voluntary participation and coerced compliance vanishes.

The UN report recommends that “explicit prohibitions of unsolicited exposure to neuromodulation procedures should be considered” (para. 29) and that “coercive uses of neurotechnologies to augment resilience on the battlefield (i.e. to reduce emotions, empathy or the need for sleep) should be prohibited” (para. 29). These recommendations apply equally to the civilian sphere: no worker should be required to submit to neural monitoring as a condition of employment; no student should be penalized for refusing cognitive surveillance; no citizen should have their neural data used to determine their access to services, insurance, or social participation.

The slippery slope is not a hypothetical. It is the path already being paved by the convergence of commercial interests, weak regulation, and the gradual normalization of neural surveillance. The UN report’s warning is clear: “without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations” (para. 13). The choice is not between neurotechnology and no neurotechnology, but between a future where consent is meaningful and one where it is a fiction that masks coercion.

PART VIII

The Benevolent Future — Neurotechnology for Good

The preceding chapters have mapped a landscape of risk: coercion, surveillance, manipulation, and the erosion of mental privacy. But to focus only on the dangers would be to ignore the transformative promise that neurotechnology holds. The UN report (A/HRC/57/61) is unequivocal: "Promising developments in scientific research demonstrate how those technologies may contribute to advancing applied neurosciences. Better understanding the functioning of the nervous system and human brain may provide new tools for the diagnosis, prevention and treatment of neurological and psychiatric diseases" (para. 1).

This is the benevolent face of neurotechnology—a future where brain-computer interfaces restore mobility to the paralyzed, where closed-loop deep brain stimulation halts Parkinson's tremors before they begin, where real-time mood monitoring prevents suicide attempts, and where children with learning disabilities receive personalized neurofeedback that rewires their neural pathways for success. The UN report notes that "assistive devices hold great potential to improve people's mobility and autonomy" (para. 1). For millions living with locked-in syndrome, spinal cord injury, or degenerative neurological conditions, neurotechnology is not a threat—it is a lifeline.

Medical Miracles Already Here

Invasive brain stimulators have been implanted worldwide for decades, treating Parkinson's disease, epilepsy, and chronic pain. The UN report confirms that "neurotechnologies have for long been used in the medical domain and are researched to treat neurological disorders,

such as Parkinson's disease, epilepsy and chronic pain, and hold promises in the treatment of depression" (para. 4). These are not speculative technologies; they are established treatments with proven efficacy.

The convergence of neurotechnology with artificial intelligence has accelerated these capabilities. The report explains that "with the support of artificial intelligence and machine learning, allowing for the processing and analysis of large sets of neurodata, the opportunities and uses of neurotechnologies are multiplying" (para. 7). Closed-loop systems—where devices both read and write neural signals in real-time—are enabling treatments that adapt autonomously to the patient's changing brain state.

Beyond Medicine: Education, Well-being, and Human Potential

The UN report acknowledges that "different applications for direct-consumer use in education, leisure and well-being are already on the market" (para. 1). Non-invasive neurotechnologies—wearable headsets, EEG headbands, and biosensor-integrated earbuds—are being used to "improve cognitive and sporting performance or to aid concentration, relaxation or sleep" (para. 10). Users can track their levels of concentration and check results on their phones, potentially revolutionizing how we understand and optimize our own cognitive states.

In education, neurotechnology could enable personalized learning experiences that adapt in real-time to a student's attention levels, cognitive load, and emotional state. The report notes that these technologies "may revolutionize the way in which people communicate or live" (para. 1). Imagine a classroom where every child receives instruction calibrated to their unique neural processing patterns—where struggling readers receive targeted neurofeedback that strengthens the neural circuits for literacy, and where gifted students are challenged at precisely the right level to maintain engagement without frustration.

The Promise and Its Conditions

The UN report is careful to distinguish between potential and reality. It notes that "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge" and that "research on potential side-effects, long-term impacts and their possible irreversible character is lacking" (para. 10). The benevolent future is not guaranteed; it must be built on a foundation of rigorous science, ethical development, and human rights compliance.

The report frames the challenge clearly: "Neurotechnologies provide a growing number of opportunities, which are to be explored, but without obscuring the profound ethical and human rights implications at stake" (para. 2). The path forward requires "developing specific policies, legislation and instruments" and "reinforcing the institutional structure to ensure that prevention and mitigation policies are effectively implemented" (para. 15).

The benevolent future of neurotechnology is not a separate path from the protective framework outlined in this guide—it is the destination that proper safeguards make possible. The UN report reminds us that "when safe, effective, secure and human-rights compliant

neurotechnology products do exist, access becomes a key element of the right to health" (para. 32). The goal is not to halt neurotechnology, but to steer it toward human flourishing, ensuring that its benefits are distributed equitably and its risks are contained by robust legal and ethical guardrails.

CHAPTER 47

Medical Miracles — Treating Parkinson's, Epilepsy, Depression, and Paralysis

Before the alarms, before the warnings about cognitive liberty and neural surveillance, there is this: the hospital room. The patient whose hand trembles uncontrollably, whose body seizes without warning, whose mind is trapped in a body that will not obey. For them, neurotechnology is not a threat to be feared but a miracle to be hoped for.

The United Nations Human Rights Council Advisory Committee, in its landmark 2024 report (A/HRC/57/61), acknowledges this directly. Neurotechnologies, the report states, "have for long been used in the medical domain and are researched to treat neurological disorders, such as Parkinson's disease, epilepsy and chronic pain, and hold promises in the treatment of depression" (para. 4). These are not speculative applications; they are the reality of modern medicine, already transforming lives around the world.

Deep Brain Stimulation: Rewiring the Diseased Brain

The most established of these medical miracles is deep brain stimulation (DBS). Invasive brain stimulators, the report notes, "have been in use for decades and are being implanted worldwide for the treatment of neurological conditions" (para. 11). The procedure is precisely what it sounds like: electrodes are surgically implanted into specific regions of the brain, connected to a pacemaker-like device placed under the skin of the chest. Electrical pulses modulate abnormal neural activity, effectively restoring function to circuits that have gone awry.

For Parkinson's disease patients, DBS can mean the difference between a life of freezing, tremors, and medication-induced dyskinesia, and a life of relative normalcy—the ability to walk, to speak, to feed oneself. For epilepsy patients, responsive neurostimulation systems can detect the onset of a seizure and deliver counteracting stimulation before the patient even knows what is happening. These are closed-loop systems, as the report describes them, where "both functions, that is reading and writing, are capable of adapting themselves autonomously" (para. 8, footnote 5). The device reads the brain's electrical activity, recognizes the signature of an impending seizure, and writes a corrective signal back into the neural circuitry—all without conscious thought from the patient.

Brain-Computer Interfaces: Restoring Movement and Communication

Perhaps the most dramatic medical applications involve brain-computer interfaces (BCIs) for paralysis. The UN report defines BCIs as "devices and methods that establish a direct connection and/or communication between the brain and external devices (i.e. phone or computer), allowing the possibility of exchanging information between the two" (para. 8). For a person with locked-in syndrome—fully conscious but unable to move a single muscle—a BCI can be a lifeline to the world.

The report describes how BCI technology works: “It involves collecting the brain’s neural activity using recording devices, decoding these activities to extract information about the individual and using this information to output corresponding commands to control external devices, thereby creating an interactive closed loop” (para. 8). A patient thinks “move cursor left,” and the cursor moves left. A patient thinks “I am thirsty,” and a speech synthesizer speaks the words. The technology is not science fiction; it is being tested in human trials today.

The report specifically notes that a company “has recently developed a secure interface for communication ‘with the power of thought’” and “is conducting large-scale trials of this technology, which can be implanted into the brain through blood vessels” (para. 11). This is Synchron’s Stentrode, a device that is threaded through the blood vessels to the brain’s motor cortex, avoiding the need for open-brain surgery. It is a remarkable example of how medical neurotechnology is becoming less invasive, more accessible, and more powerful.

The Treatment-Resistant Depression Frontier

Beyond movement disorders and paralysis, neurotechnology is opening new frontiers in mental health. The UN report notes that DBS “holds promises in the treatment of depression” (para. 4). For the millions of people worldwide who do not respond to medication or therapy—treatment-resistant depression—this represents a potential revolution. Closed-loop systems can now monitor mood-related neural patterns and deliver stimulation precisely when needed, potentially preventing depressive episodes before they fully manifest.

The report frames this within the right to health: “The development of neurotechnology applications holds great prospects in protecting mental health… and granting personal autonomy” (para. 31). When a patient’s own brain chemistry is working against them, neurotechnology can be the tool that restores balance—not by chemical intervention alone, but by directly modulating the neural circuits that underpin mood, cognition, and emotion.

The Critical Caveat: Safety and Evidence

The UN report is careful to distinguish between what is proven and what is promised. It notes that in the medical realm, neurotechnologies “are generally regulated by established frameworks and processes” (para. 12). This is not the Wild West of consumer neurotech; medical devices must undergo rigorous clinical trials, safety testing, and regulatory approval before they can be implanted in patients.

Yet even in medicine, the report warns, “the potential long-term effects of using neurotechnologies are largely unknown” (para. 31). Patients who have undergone DBS have reported “feeling a changed sense of agency and identity” (para. 27, footnote 36). The very technology that restores function can also alter personality, blurring the line between treatment and transformation. Ensuring “psychological continuity” becomes a medical and ethical imperative.

The Human Rights Framework for Medical Neurotechnology

The UN report grounds these medical applications in human rights law. Patients have a "right to choose or refuse the treatment they want with the full knowledge of the risks and benefits" (para. 28). This is the principle of informed consent, and it applies with special force to neurotechnology, where the stakes include not just physical health but mental integrity, personal identity, and cognitive liberty.

For persons with disabilities, the report emphasizes that the Convention on the Rights of Persons with Disabilities "provides a protective framework that cannot be departed from" (para. 39). Neurotechnology must be developed "with meaningful engagement and participation of persons with disabilities" (para. 38), not imposed upon them. The goal is not to "fix" disabled brains but to restore function and autonomy in a way that respects the dignity and self-determination of the individual.

When safe, effective, and human-rights compliant neurotechnology products do exist, the report insists, "access becomes a key element of the right to health" (para. 32). States must "grant access, without discrimination, in accordance with national health service regulations and the principle of progressive realization" (para. 32). The medical miracle is not truly a miracle if it is only available to the wealthy.

The Bottom Line

Neurotechnology is already saving lives, restoring function, and alleviating suffering. The UN report recognizes this potential while insisting that it must be pursued within a human rights framework. The medical miracles of DBS, closed-loop systems, and BCIs for paralysis are not arguments against regulation—they are arguments for regulation that is smart, targeted, and protective. The goal is to maximize the therapeutic benefit while minimizing the risk of harm, ensuring that the technology serves the patient, not the other way around.

CHAPTER 48

Cognitive Enhancement — Smarter, Faster, Better?

The same neurotechnologies that restore lost function can also augment healthy brains. This is the promise—and the peril—of cognitive enhancement: the use of neurotechnology to make healthy people smarter, faster, more focused, and more capable than they would otherwise be. The UN report addresses this frontier directly, noting that devices are being developed for "the 'cognitive enhancement, pleasure or leisure of healthy people'" (para. 11). The question is no longer whether such enhancement is possible, but who gets access, under what conditions, and at what cost to equity and human rights.

The Enhancement Landscape

Cognitive enhancement takes many forms. At the low-tech end, nootropics—so-called "smart drugs"—have been used for years by students and professionals seeking a cognitive edge. But neurotechnology offers something far more direct: the ability to stimulate, modulate, or augment specific brain functions through electrical or magnetic means.

Non-invasive devices are already marketed directly to consumers for cognitive enhancement. As the UN report notes, "the most popular wearables are used to improve cognitive and sporting performance or to aid concentration, relaxation or sleep" (para. 10). These devices, often in the form of headsets or earbuds, claim to track and enhance mental states. However, the report cautions that "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge" (para. 10). Moreover, "research on potential side-effects, long-term impacts and their possible irreversible character is lacking" (para. 10).

More invasive approaches are on the horizon. The report references companies that "already advertise 'cosmetically' invisible implants that may allow users to control computers or mobile devices from any location" and that are "working on the 'cognitive enhancement, pleasure or leisure of healthy people'" (para. 11). Neuralink, for example, "claims to have developed a brain implant that will allow a connection to the Internet at the biological level, which will allow people to communicate wirelessly with anyone with a similar implant and computer setup" (para. 11, footnote 9). These are not speculative futures; they are technologies currently in human trials.

The Human Rights Challenge

The UN report makes clear that cognitive enhancement for healthy persons raises fundamentally different questions than therapeutic use. In the medical sphere, the risk-benefit assessment is guided by the principle of *primum non nocere*—first, do no harm. But for consumer products promoted for cognitive enhancement, "the risk-benefit assessment may not be the same" (para. 31). The report warns that "in the absence of a medical need and without being supported by extensive trials and scientific evidence, devices may become

harmful for mental health and their extensive use may lead to public health problems" (para. 31).

The report also identifies a critical gap in regulation. While medical neurotechnologies are "generally regulated by established frameworks and processes, the consumer sector remains highly deregulated and is characterized by rapid expansion" (para. 12). This deregulation creates a permissive environment for enhancement technologies to reach the market without adequate safety testing or human rights safeguards.

The Access Question: Who Gets Enhanced?

Perhaps the most profound challenge posed by cognitive enhancement is the question of equitable access. The UN report notes that "when safe, effective, secure and human-rights compliant neurotechnology products do exist, access becomes a key element of the right to health" (para. 32). But the report also acknowledges that "universal access to advanced technology depends, however, on needs assessments and budgetary decisions" (para. 32). In practice, this means that enhancement technologies—unlike therapeutic interventions—are likely to be available first to those who can afford them.

The report flags a crucial distinction: "Although access to 'neuroenhancement' is often framed in terms of 'equality' and 'democratization,' a human rights approach requires granting priority access to those in need" (para. 32, footnote 43). This is a direct challenge to the narrative that enhancement technologies will inevitably level the playing field. In reality, without deliberate policy interventions, cognitive enhancement risks widening existing inequalities, creating a new class of cognitively augmented elites who can afford to boost their own capabilities while the unenhanced majority falls further behind.

The Coercion Risk

The UN report also warns against coercive uses of enhancement technology, particularly in contexts characterized by power asymmetries. "Coercive uses of neurotechnologies to augment resilience on the battlefield (i.e. to reduce emotions, empathy or the need for sleep) should be prohibited, while non-coercive uses may require limitations" (para. 29). This is not a hypothetical concern. Military programs like DARPA's N3 initiative are actively developing non-surgical neurotechnologies for able-bodied soldiers, raising questions about whether participation is truly voluntary in a hierarchical, command-driven environment.

The same coercion risks exist in civilian contexts. Employers may pressure workers to use enhancement devices to increase productivity. Schools may encourage or require students to use attention-monitoring headsets. The line between voluntary enhancement and coerced optimization is dangerously thin, and the UN report's insistence on "meaningful and effective free and informed consent" (para. 28) applies with full force to enhancement contexts.

The Bottom Line

Cognitive enhancement is not a distant possibility—it is already happening, in a regulatory vacuum, driven by commercial interests and military imperatives. The UN report calls for a human rights framework that prioritizes safety, equity, and consent. Enhancement technologies must be held to the same standards of evidence and ethical oversight as therapeutic interventions. And society must confront the fundamental question: do we want a world where cognitive ability becomes a commodity that can be bought and sold, or do we want to preserve the principle that human potential should not be limited by the size of one's bank account?

CHAPTER 49

Neuroeducation — Learning at the Speed of Thought

The same neurotechnologies that raise profound human rights concerns in military and commercial contexts also hold transformative potential for one of society's most fundamental institutions: education. The United Nations Human Rights Council Advisory Committee, in its landmark 2024 report (A/HRC/57/61), explicitly identifies "different applications for direct-consumer use in education" as already being on the market (para. 1). The question is no longer whether neurotechnology will enter the classroom, but under what conditions, with what safeguards, and for whose benefit.

Personalized Learning at Neural Resolution

The core promise of neuroeducation is the ability to tailor instruction to the individual brain's real-time state. Non-invasive EEG headsets, already commercialized for consumer use, can wirelessly track a student's levels of concentration and display results on a phone (para. 10). In theory, this allows teachers to know precisely when a student is engaged, confused, distracted, or ready to move to the next concept—without relying on the imperfect signals of raised hands or glazed eyes.

The UN report notes that such devices are "used for a variety of medical purposes, but the same devices are also being commercialized for different applications in the area of well-being, in the workplace or as an educational tool" (para. 10). This dual-use character means that technologies developed for clinical settings—monitoring attention deficits in ADHD patients, for example—can be repurposed for classroom-wide deployment with healthy students.

The Attention Monitoring Dilemma

The most immediate application of neurotechnology in education is attention monitoring. Students wearing non-invasive headsets can have their neural activity continuously measured, with algorithms determining whether they are focusing on the lesson or daydreaming. The data can be aggregated for the teacher, shared with parents, or even used to adjust the difficulty of material in real time.

But the UN report sounds a clear warning: "The effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge. Moreover, research on potential side-effects, long-term impacts and their possible irreversible character is lacking" (para. 10). This is not a minor caveat—it is a fundamental challenge to the entire enterprise of neuroeducation as currently practiced. Schools may be adopting technologies whose benefits are unproven and whose risks are unknown.

The report further notes that "neurotechnologies can be used to interfere and manipulate individuals" and that "neurogaming devices can also drive compulsive or addictive use,

particularly when neural reward circuits target altering consumption behaviours" (para. 21). If educational neurotechnology is designed to reward certain brain states with positive feedback, the line between healthy engagement and neural conditioning becomes dangerously thin.

The Consent Problem in the Classroom

Informed consent is a cornerstone of human rights protection for neurotechnology. The UN report insists on "meaningful and effective free and informed consent" (para. 28). But how does this apply to children in a compulsory education setting?

The report acknowledges that "individuals may not fully understand the implications of using neurotechnologies" and "may unknowingly surrender control over their cognitive processes" (para. 30). For children, whose cognitive faculties are still developing and whose capacity for informed consent is legally limited, this concern is magnified. The report specifically warns that "negative impacts on mental health of children, older persons and persons with disabilities should be considered together with the opportunities" (para. 33).

A student who refuses to wear an attention-monitoring headset may be labeled as uncooperative or may miss out on personalized learning benefits that their peers receive. The coercive dynamic is built into the structure of the classroom, where teachers hold authority and students have limited power to opt out.

Accelerated Skill Acquisition and Its Limits

Beyond attention monitoring, neurotechnology promises accelerated skill acquisition through techniques like neurofeedback and transcranial stimulation. The idea is that by reinforcing optimal neural patterns during learning—whether for language acquisition, mathematical reasoning, or musical performance—students can achieve in months what previously took years.

The UN report is cautiously optimistic about such applications, noting that neurotechnologies "hold great potential to improve people's mobility and autonomy" and "may revolutionize the way in which people communicate or live" (para. 1). But it immediately adds a crucial caveat: "But to do so, direct access to a person's brain activity has to be granted to external devices" (para. 1). The very mechanism that enables accelerated learning also creates a gateway for interference, surveillance, and manipulation.

The Neuronal Divide in Education

The UN report warns that "groups experiencing structural discrimination and marginalized people or those in positions of socioeconomic disadvantage are the most at risk of misuse of neurotechnologies during the trial phase" (para. 37). In education, this means that the poorest students—those in underfunded schools with the least parental advocacy—may be the first subjected to experimental neurotechnology programs with inadequate safeguards.

Conversely, once neurotechnologies are proven safe and effective, the report notes that "the members of those groups are usually excluded from or discriminated against in accessing

those technologies, which may not be affordable" (para. 37). The result is a double injustice: the vulnerable are used as test subjects during the experimental phase, then priced out of the benefits once the technology matures.

This creates the risk of a "neuronal divide" in education, where wealthy students access cutting-edge cognitive enhancement tools that boost their learning efficiency, while poorer students receive only the surveillance and behavioral modification aspects of the technology—or none at all.

The Regulatory Vacuum

The UN report paints a stark picture of the current regulatory landscape: "In most countries, applicable regulations are unclear, weak or non-existent. Existing loopholes in regulations, lack of technical expertise and capacity and the absence of adequate oversight bodies are factors that will certainly be exploited by large companies seeking profits" (para. 13).

For neuroeducation, this means that school districts and education ministries are making decisions about technology adoption without clear guidance from human rights frameworks. Companies selling "learning enhancement" headsets face minimal scrutiny regarding their data practices, the accuracy of their algorithms, or the long-term effects on developing brains.

The report calls for action: "Where areas not sufficiently covered by existing standards do emerge, action is needed to prevent violations, misuses and abuses. Developing context-specific standards may be a necessary measure to better define the scope of rights being engaged" (para. 14). Neuroeducation is precisely such an area—a domain where the promise of accelerated learning must be balanced against the fundamental rights of children to privacy, mental integrity, and freedom from manipulation.

CHAPTER 50

Mental Health — New Frontiers in Treatment

Among the most promising applications of neurotechnology is its potential to transform mental health care. As the 2024 UN report notes, "Better understanding the functioning of the nervous system and human brain may provide new tools for the diagnosis, prevention and treatment of neurological and psychiatric diseases" (A/HRC/57/61, para. 1). For the millions suffering from depression, PTSD, anxiety, and other mental health conditions, neurotechnology offers hope where traditional treatments have failed.

Real-Time Mood Monitoring

One emerging frontier is the ability to monitor mental states in real time. Non-invasive wearables—headbands, earbuds with biosensing electrodes, and wristbands—can now track brain activity patterns associated with different emotional states. The UN report observes that "users can wirelessly track their levels of concentration and check their results on their phone" (para. 10). While these consumer devices are often marketed for wellness rather than clinical treatment, the underlying technology points toward a future where mood disorders can be detected and addressed before they escalate.

The report cautions, however, that "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge. Moreover, research on potential side-effects, long-term impacts and their possible irreversible character is lacking" (para. 10). This gap between commercial promise and clinical validation creates both opportunity and risk for mental health applications.

Closed-Loop Depression Treatment

Perhaps the most revolutionary development is the closed-loop system—a neurotechnology that can both "read" and "write" neural activity. The UN report explains that "closed-loop systems result from the fact that both functions, that is reading and writing, are capable of adapting themselves autonomously" (para. 7, footnote 5). In the context of depression treatment, this means a device can continuously monitor a patient's neural state and deliver precisely calibrated stimulation when it detects the onset of depressive symptoms.

Deep brain stimulation (DBS) has already been used for decades to treat neurological conditions like Parkinson's disease. The UN report confirms that "invasive brain stimulators have been in use for decades and are being implanted worldwide for the treatment of neurological conditions" (para. 11). The next generation of these devices, enhanced by AI-powered closed-loop algorithms, could offer on-demand relief for treatment-resistant depression—effectively acting as a "pacemaker for the brain."

The potential is immense, but so are the ethical challenges. The UN report emphasizes that "patients having undergone deep brain stimulation have reported feeling a changed sense of agency and identity; thus, ensuring 'psychological continuity' may be important" (para. 27,

footnote 36). When a device can alter your mood, your sense of self, and your emotional responses, fundamental questions arise about who is controlling whom.

PTSD Therapy and Trauma Processing

Neurotechnology also offers new avenues for treating PTSD. By combining real-time neural monitoring with therapeutic interventions, clinicians could potentially guide patients through trauma processing while ensuring they remain within a therapeutic window—not so overwhelmed that they retraumatize, but engaged enough to process the experience.

The UN report frames these developments within the right to health, noting that "the development of neurotechnology applications holds great prospects in protecting mental health and granting personal autonomy" (para. 31). However, it immediately warns that "the right to health may also be threatened through the use and misuse of neurotechnologies" (para. 31). This dual-edged nature is particularly acute in mental health, where the line between therapy and manipulation can be thin.

The Consent Challenge

A critical concern specific to mental health applications is informed consent. The UN report highlights that "in the case of consumer-oriented neurotechnology devices, it is hard to see how users can provide informed consent when their potential long-term effects are largely unknown" (para. 30). For patients with depression, PTSD, or other conditions that impair decision-making capacity, this challenge is magnified.

The report calls for reinforced standards: "Ensuring that decisions in neurotechnology interventions and clinical trials are genuinely autonomous is vital. Human rights standards and limitations in relation to medical practices and coercive health treatments, particularly in psychiatry, also apply to neurointerventions" (para. 28). States must guarantee people's "right to choose or refuse the treatment they want with the full knowledge of the risks and benefits" (para. 28, quoting Committee on Economic, Social and Cultural Rights).

Access and Equity

Finally, the UN report raises the question of access. When safe and effective neurotechnology treatments do exist, "States should then grant access, without discrimination, in accordance with national health service regulations and the principle of progressive realization" (para. 32). But the reality is that "cutting-edge treatments, such as implants controlled by brain-computer interface software, are often costly and require highly specialized surgeries, which may not always be covered by public health" (para. 32, footnote 44).

The result could be a two-tier system of mental health care: those who can afford advanced neurotechnology treatments, and those who cannot. The UN report warns that "persons with disabilities, in particular, must be granted access to affordable assistive technologies without discrimination" (para. 32, citing Convention on the Rights of Persons with Disabilities, art. 20). But without deliberate policy interventions, the market will likely deliver these life-changing treatments first to the wealthy, deepening existing health inequities.

PART IX

AI and Neurotechnology — The Convergent Revolution

The convergence of artificial intelligence and neurotechnology represents perhaps the most significant inflection point in the history of human rights. As the United Nations Human Rights Council Advisory Committee observed in its landmark 2024 report, "With the support of artificial intelligence and machine learning, allowing for the processing and analysis of large sets of neurodata, the opportunities and uses of neurotechnologies are multiplying." This is not merely a matter of faster processing or bigger datasets. AI fundamentally transforms what neurotechnology can do—and what it can do *to* us.

The AI-Neuro Symbiosis

Neurotechnology without AI is like a car without an engine: it has form but no power to move. The raw neural signals captured by electroencephalography (EEG) or functional magnetic resonance imaging (fMRI) are noisy, complex, and difficult to interpret. AI algorithms—particularly deep learning models—excel at extracting meaningful patterns from this chaos. They can decode neural activity into specific mental states: memories, emotions, intentions, even inner speech.

The UN report notes that this capability is still imperfect: "Despite their rapid progress, most neurotechnologies are still being tested and do not provide accurate results in the decoding of people's thoughts, which means that 'mind-reading' or 'mind-steering' are not still achievable." But the trajectory is clear. Each year brings improvements in signal processing, pattern recognition, and predictive accuracy. The gap between current capability and science fiction is narrowing.

The Data Engine

AI does not simply interpret neurodata—it also drives its collection. The UN report explains that neurotechnologies "collect, analyse and process large personal datasets of a highly sensitive nature." These datasets are not static; they grow with every use, every calibration, every optimization. AI systems require vast amounts of training data to function effectively, and the companies developing neurotechnology products have every incentive to collect as much as possible.

This creates a feedback loop: more data enables better AI, which enables more capable neurotechnology, which collects more data. The UN report warns that "consumer neurotechnology companies are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future." The data collected today may reveal far more tomorrow, as AI decoding algorithms improve.

Algorithmic Bias and Neurodiscrimination

The convergence of AI and neurotechnology also amplifies existing risks of algorithmic bias. The UN report notes that "potential misuses of such information are a major concern and issues resulting from the convergence with artificial intelligence, such as algorithmic bias and lack of traceability and transparency, will also be exacerbated through their use."

If AI training data for neurotechnology is drawn primarily from certain populations—for example, young, healthy, affluent individuals from developed countries—the resulting systems may perform poorly for others. A brain-computer interface trained on the neural patterns of one demographic may misinterpret the signals of another, leading to errors in diagnosis, communication, or control. The UN report emphasizes that "groups experiencing structural discrimination and marginalized people or those in positions of socioeconomic disadvantage are the most at risk of misuse of neurotechnologies."

The Autonomy Challenge

Perhaps the most profound implication of the AI-neuro convergence concerns human autonomy. When AI systems make decisions based on neural data, the question of who is truly in control becomes urgent. The UN report frames this as a challenge to "the foundational principle of human dignity," which requires assessing "whether autonomy may be neglected by altering the foundation of their mental processes and states."

Closed-loop neurotechnology systems—those that both read and write neural activity—represent the frontier of this challenge. The UN report describes how "closed-loop systems result from the fact that both functions, that is reading and writing, are capable of adapting themselves autonomously." These systems can adjust their stimulation parameters in real time, without human intervention, based on AI analysis of the user's neural state. The user may not even be aware that their brain activity is being modulated.

A Rapidly Growing Industry

The market forces driving this convergence are immense. The UN report notes that private investment in neurotechnology "increased from $331 million to $7.3 billion in just 10 years." Nearly half of all worldwide patent applications in the field are registered in the United States. Major technology companies are racing to bring neurotechnology products to market, and the pressure will only increase.

Yet this rapid growth is occurring in a regulatory vacuum. The UN report warns that "in contrast to the medical development of neurotechnologies, which are generally regulated by established frameworks and processes, the consumer sector remains highly deregulated and is characterized by rapid expansion." The AI systems that power these consumer devices are subject to even less oversight.

The Urgent Need for Guardrails

The UN report concludes that "the risk is that, without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations." The convergence of AI and neurotechnology demands a corresponding convergence of regulatory frameworks. Human rights impact assessments, transparency requirements, and independent oversight mechanisms must be developed for both technologies—and for their combination.

As the UN report states, "When dealing with highly disruptive technologies anticipationis key. Although human rights violations may happen as a consequence of actions and misuses, they are also the result of lack of preparedness."

The convergence of AI and neurotechnology is not a future possibility—it is a present reality, accelerating daily. The question is whether our legal and ethical frameworks can keep pace.

CHAPTER 51

Human-Machine Teaming — The Centaur Model

The most transformative vision for neurotechnology is not one in which machines replace humans, nor one in which humans remain separate from machines. It is the "centaur" model: a seamless partnership in which human and artificial intelligence collaborate, each contributing its distinctive strengths. As former Deputy Secretary of Defense Robert Work described this emerging paradigm, "The coin of the realm is going to be learning machines and human-machine collaborations, which allows machines to allow humans to make better decisions; assisted human operations, which means bringing the power of the network to the individual; human-machine combat teaming; and the autonomous network."

The centaur metaphor is apt. In Greek mythology, the centaur combined human intelligence with equine strength and speed. In the neurotechnology age, the centaur combines human creativity, judgment, and ethical reasoning with machine speed, memory, and pattern recognition. The bridge between them is the brain-computer interface (BCI), which enables bidirectional communication between the human brain and external devices—allowing humans to think with machines, not merely through them.

The Cognitive Partnership

Current human-machine interaction is constrained by physical interfaces: keyboards, screens, voice commands. These bottlenecks limit the speed and bandwidth of communication. BCI bypasses these constraints. As the RAND Corporation's 2020 assessment of BCI military applications explains, "BCI provides a mechanism for blending human strengths and computer strengths, and much of the ongoing work strives to link these two sets of capabilities and yield synergistic advantages."

This goes beyond simple control. The concept of "cortically coupled AI" represents a heightened level of teaming, where "a cortically coupled computer system opportunistically senses the brain state, capturing a user's implicit or explicit computation, and then communicates this information to a traditional computer system via a neural interface." Rather than typing commands or speaking instructions, the human thinks them—and the machine responds. This information can then help train AI systems, creating a feedback loop of mutual enhancement.

Elon Musk, founder of Neuralink, framed the ambition directly: "Some high bandwidth interface to the brain will be something that helps achieve a symbiosis between human and machine intelligence and maybe solves the control problem and the usefulness problem."

Operational Applications

The RAND study tested the centaur model through a table-top exercise involving tactical urban operations, using a projected "BCI toolbox" of capabilities. Participants identified three areas where human-machine teaming through BCI offered transformative potential.

First, **enhanced speed of communication**. On a future battlefield where information flows at machine speed, human reaction times become the critical vulnerability. BCI can compress the observe-orient-decide-act (OODA) loop by allowing thoughts to become actions directly, without the delay of physical movement. A quadriplegic woman operating an F-35 simulator with her thoughts—a capability already demonstrated by DARPA, the Applied Physics Laboratory, and the University of Pittsburgh—illustrates the principle, even if the application differs.

Second, **improved common situational awareness**. BCI could aggregate the cognitive states of an entire squad, allowing commanders to monitor cognitive workload, stress levels, and attention in real time. The Army Research Laboratory is already using 3D-printed helmets with integrated EEG sensors to monitor brain activity. The Air Force is pursuing camera-based cognitive monitoring systems built into pilots' helmets to "adapt displays based on the pilot's unique physical and mental conditions."

Third, **simultaneous control of multiple platforms**. Researchers at the Human-Oriented Robotics and Control laboratory, with funding from DARPA and the U.S. Army, have enabled a user to control a swarm of drones using BCI. The lab's researchers suggest the technology could be practical for military use within five to ten years, with applications including delivery of medical help, search and rescue, and exploration.

Brain-to-Brain Communication

The most speculative—and potentially most transformative—application of the centaur model is direct brain-to-brain communication. Researchers at the University of Washington, with funding from the Army Research Laboratory, conducted a pilot study using a noninvasive system that reads basic brain signals via EEG, transmits them over the internet, and transfers motor responses to a second user using transcranial magnetic stimulation. The signals represented basic actions in a simple video game—move left or right—but the implications extend far beyond gameplay.

The researchers extended this work to involve groups of three individuals: two sending information, one receiving, all collectively playing a Tetris-like game. Earlier experiments had transferred signals between rats and from a human to a rat. The ability to send even basic thoughts across the internet "inherently presents many opportunities and many risks concerning security and ethics," the RAND study notes.

The Trust Deficit

Despite the promise, the centaur model faces a critical barrier: trust. The RAND study found that "cultural barriers to BCI, particularly among infantry service members, are likely to be high." Soldiers may resist having their brains connected to machines, fearing loss of autonomy, vulnerability to hacking, or simply the unnaturalness of the interface.

The study recommends mitigating this trust deficit through "heavy vetting and testing in noncombat scenarios, introduction to service members that already rely on machine

technologies, and an initial focus on noninvasive measures and medical applications." Trust must be earned, not imposed.

The Institutional Challenge

Human-machine teaming through BCI will require not just technological advances but institutional innovations. As the RAND study concludes, "As the U.S. government prepares to incorporate BCI technologies into future military capabilities, it will require institutional innovations to address new ethical and policy issues at each stage of the process, from R&D to operational application to veteran care."

The centaur model promises enhanced decision-making, creativity, and problem-solving by combining human and machine intelligence. But it also raises profound questions: Who is responsible when a human-machine team makes a mistake? What happens when the machine's recommendation conflicts with the human's judgment? And how do we ensure that the centaur serves human purposes, rather than the reverse?

CHAPTER 52

AI as the Brain's Operating System

The convergence of artificial intelligence and neurotechnology is not merely incremental—it is transformative. As the 2024 UN Human Rights Council report (A/HRC/57/61) makes clear, it is "with the support of artificial intelligence and machine learning, allowing for the processing and analysis of large sets of neurodata, [that] the opportunities and uses of neurotechnologies are multiplying." Without AI, most neurotechnologies would remain crude measurement tools. With AI, they become adaptive systems capable of decoding, predicting, and modulating mental states in real time.

The Decoding Layer

At its most fundamental level, AI enables the translation of raw neural signals into meaningful information. The human brain produces electrical activity that is noisy, complex, and highly individualized. An electroencephalogram (EEG) recording, for example, captures millions of neural firings per second, mixed with muscle artifacts, environmental interference, and signal drift. Without AI, extracting a user's intention—say, to move a cursor left—from this cacophony is nearly impossible.

Machine learning algorithms, particularly deep neural networks, have revolutionized this process. The UN report notes that brain-computer interfaces (BCIs) function by "collecting the brain's neural activity using recording devices, decoding these activities to extract information about the individual and using this information to output corresponding commands to control external devices, thereby creating an interactive closed loop." The decoding step is where AI is essential. Algorithms are trained on vast datasets of neural recordings paired with known mental states or intended actions. Once trained, they can recognize patterns in real time, translating thought into action with increasing accuracy.

This capability extends beyond simple motor commands. The UN report observes that "neuroimaging is, however, increasingly explored and can be employed to make inferences about individuals on a large spectrum of mental states, including memories, semantic knowledge, emotions, dreams, inner speech and intentions." AI systems can now predict what word a person is about to speak, what image they are visualizing, or even what memory they are recalling, all from patterns of brain activity.

The Predictive Layer

Beyond decoding current mental states, AI enables the prediction of future states. This is where the technology becomes most powerful—and most concerning. By analyzing an individual's neural patterns over time, AI systems can learn to anticipate emotional shifts, cognitive fatigue, or even behavioral intentions before the person is consciously aware of them.

The UN report highlights that "with the support of artificial intelligence, neurotechnologies can elaborate sophisticated individualized psychological profiles and use them to predict individuals' behaviour or intentions or to infer their mental characteristics or predispositions." This predictive capability has profound implications for mental privacy. A system that can predict your emotional state five minutes from now, or your likelihood of making a particular decision, has access to information that even you may not yet possess.

Consider a consumer EEG headband marketed for focus enhancement. The AI behind it learns your concentration patterns throughout the day, identifying when you are most productive and when you are prone to distraction. Over time, it can predict your optimal work schedule, recommend when to take breaks, and even suggest when to consume caffeine. This seems benign—even helpful. But the same predictive algorithms could be used by employers to anticipate which employees are likely to quit, by insurers to predict mental health episodes, or by governments to identify potential dissidents.

The Optimization Layer

The most sophisticated application of AI in neurotechnology is the closed-loop system, where AI not only reads neural signals but also writes back to the brain in real time. The UN report describes these systems as those where "both functions, that is reading and writing, are capable of adapting themselves autonomously." This is the brain's operating system in its fullest expression.

In medical contexts, closed-loop deep brain stimulation (DBS) systems use AI to monitor neural activity and adjust stimulation parameters automatically. A patient with Parkinson's disease, for example, might have a DBS implant that detects the onset of tremor and delivers precisely calibrated electrical pulses to suppress it—all without the patient's conscious intervention. The AI learns the patient's unique neural signature for tremor and optimizes stimulation over time, becoming more effective with each cycle.

The UN report notes that these systems are "nurtured by 'neurodata', which are needed for their own functioning, calibration and optimization." This creates a feedback loop: the more data the system collects, the better it performs; the better it performs, the more users rely on it; the more they rely on it, the more data it collects. This cycle drives rapid improvement but also raises questions about dependency and control. When an AI system is optimizing your brain function, who is really in charge?

The Risks of Algorithmic Opacity

The integration of AI into neurotechnology introduces a critical vulnerability: algorithmic opacity. The UN report warns that "issues resulting from the convergence with artificial intelligence, such as algorithmic bias and lack of traceability and transparency, will also be exacerbated through their use." When an AI system makes a decision about your neural data—whether to adjust stimulation, classify a mental state, or predict a behavior—the reasoning behind that decision may be inscrutable.

Deep neural networks, the most powerful AI models, are often described as "black boxes." Their internal workings are so complex that even their creators cannot always explain why they produce particular outputs. In the context of neurotechnology, this opacity is dangerous. If a BCI misclassifies a user's intention, causing a prosthetic limb to move incorrectly, the consequences could be physical harm. If an AI system falsely predicts that a person is experiencing suicidal ideation, leading to involuntary psychiatric intervention, the consequences could be a violation of fundamental rights.

The UN report emphasizes that "human rights impact assessments should also be required" for neurotechnology products. But how can an impact assessment account for the behavior of an AI system that evolves over time, learning from each user's neural data? Traditional safety testing, which evaluates a device at a fixed point in time, is inadequate for adaptive systems.

The Commercial Imperative

The pressure to deploy AI-powered neurotechnology is immense. The UN report notes that "the growing role of the private sector is changing the dynamics of the industry, which is leaning towards the production of direct-to-consumer devices." Companies like Neuralink and Synchron are racing to bring brain-computer interfaces to market, and AI is their key enabling technology.

The report warns that "the possibility that, in the coming years, those products with inadequate safety measures and unclear or underestimated human rights risks may be widely commercialized is real." The combination of AI and neurotechnology creates a perfect storm: a technology that can read and write neural data, powered by algorithms that are opaque and self-optimizing, deployed in a regulatory vacuum where "applicable regulations are unclear, weak or non-existent."

The Need for AI-Specific Neuro-Rights

The convergence of AI and neurotechnology demands that neuro-rights frameworks explicitly address the role of algorithms. The right to mental privacy, for example, must include protection against AI systems that can infer mental states from neural data. The right to personal identity must guard against AI-driven modulation that could alter personality or preferences without conscious consent. The right to free will must ensure that AI recommendations do not become covert commands.

The UN report concludes that "developing context-specific standards may be a necessary measure to better define the scope of rights being engaged." As AI becomes the operating system of the human brain, those standards must be developed with urgency. The technology is already here, adapting and learning, and it will not wait for the law to catch up.

CHAPTER 53

Foundation Models for the Brain — The Coming Neural Decoder Revolution

The application of foundation-model architectures — the large-scale transformer-based systems that power modern artificial intelligence — to neural data is rapidly transforming what is technically possible. The same architectures that enable large language models to predict words can be trained to predict neural states, decode imagined speech, and reconstruct visual experience.

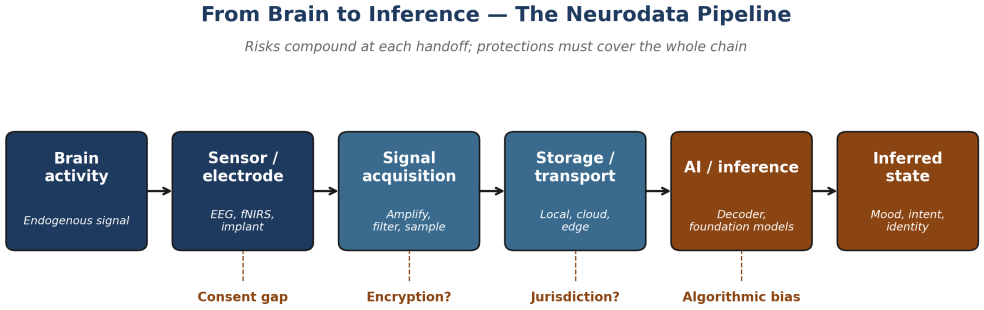


Figure 38.1. *From brain to inference — the neurodata pipeline. Risks compound at every handoff between sensor, transport, storage, and AI inference; effective protection must cover the entire chain rather than any single layer.*

From GPT to NeuroGPT

The transformer architecture, introduced in 2017 by researchers at Google, transformed natural language processing by enabling models to learn complex statistical patterns from vast amounts of text. The same architecture, applied to neural recordings, has produced startling capabilities in the past three years. Meta's 2023 demonstration of intelligible speech reconstruction from non-invasive MEG recordings, and similar work at Stanford, UCSF, and the Max Planck Institute, indicate that the application of foundation-model approaches to neural data is producing rapid capability gains.

The underlying mathematics is the same: a model trained on large datasets of neural activity, paired with the corresponding mental states or external events, learns to predict one from the other. As neural data scales — from individual laboratory recordings to industrial-scale collections from consumer devices — the predictive accuracy of these models will follow scaling curves familiar from language modelling.

The Scaling Hypothesis Applied to Neural Data

The "scaling hypothesis" in machine learning holds that model performance improves predictably with increases in three factors: model size, training data, and compute. Language

models trained with billions of parameters and trillions of tokens of text vastly outperform earlier systems. Whether the same scaling applies to neural data is an active research question, but early indications suggest that it does.

If the scaling hypothesis holds for neural data, the implications are profound. The largest current public datasets of neural recordings number in the hundreds of subjects and thousands of hours. The largest private datasets, held by companies developing consumer neurotechnology, may already exceed these by orders of magnitude. Within five years, neural datasets comparable in scale to the text datasets that trained GPT-4 may be available to well-funded developers. The decoding capabilities those datasets could enable include accurate inference of attention, emotional state, semantic content, and visual experience.

⚠ The Decoder Asymmetry

A central concern in the neuro-rights debate is the decoder asymmetry. The party that collects neural data and trains a foundation model on it acquires capabilities that the person whose brain produced the data does not possess. The data subject cannot decode their own neural signals; the data controller can. This asymmetry creates power relations that have no precedent in earlier data-protection regimes.

Cross-Subject Generalization

A historical limitation of brain decoding was that models trained on one person's neural data could not be applied to another's. Each individual's neural signatures were sufficiently idiosyncratic that decoders required per-subject calibration, limiting commercial scalability.

Foundation-model approaches are breaking this limitation. Models trained on the pooled neural data of many subjects develop representations general enough to apply to new subjects with minimal calibration. The implications for commercial deployment are substantial: a consumer device need no longer be individually trained for each user. The implications for human rights are alarming: a decoder trained on consenting research participants can in principle be applied to non-consenting targets.

The Convergence of Modalities

Foundation models excel at integrating data across modalities. The same architecture that has enabled image-text models such as CLIP, and text-image generators such as Stable Diffusion, can integrate neural data with text, image, audio, and video. The result is a class of models that can map between mental states and external representations: from neural data to predicted words spoken silently, from neural data to predicted images viewed, from neural data to predicted memories recalled.

The 2023 work by Tang and colleagues at the University of Texas demonstrated reconstruction of semantic content from fMRI of subjects passively listening to spoken narrative. The reconstructions were not literal transcriptions, but they captured semantic

content with sufficient fidelity to identify the topic, sentiment, and key entities. With larger datasets and better models, the gap between current capability and literal mind-reading will close.

What the Public Sees Versus What Industry Achieves

A consistent pattern in foundation-model development is that public academic publication lags industry capability by one to three years. The largest, best-performing models in any given period are typically held by private companies and disclosed only partially. For neural decoding, this means that public estimates of current capability are likely conservative.

Companies with the largest neural datasets — consumer device manufacturers, healthcare-system data brokers, and any state actor with surveillance capacity — may already possess decoding capabilities that have not been disclosed. The opacity of current development creates a particular risk: the public debate about neural decoding may be premised on capabilities that are already obsolete in the laboratories that matter.

Implications for Neuro-rights

The application of foundation-model approaches to neural data has three concrete implications for the neuro-rights debate:

30. **Neural data must be treated as latent information of unknown future content.** The data that is collected today, with no current decoding capability, may be decoded tomorrow with capabilities that did not exist at the time of collection. Retention policies must reflect this asymmetry.
31. **Algorithmic inferences must be protected with the same rigor as raw data.** The "neural data" definition that protects only the raw signal is insufficient. The model outputs — the inferred mental states — are the actually sensitive information.
32. **The collection of large-scale neural datasets is itself a regulatory concern.** Public policy must address not only specific uses of neural data, but the accumulation of datasets sufficient to train decoders capable of mass mental surveillance.

The Window of Opportunity

The technical capability gap between current neural decoding and the level required for mass deployment is closing rapidly. The legal and institutional response required to constrain this technology operates on a much slower timescale. Without coordinated international action in the next three to five years, the technology will outpace any meaningful regulatory framework.

The opportunity exists now to establish the legal and ethical infrastructure that will determine whether foundation-model neural decoding becomes a tool of mass surveillance or a constrained technology used only with meaningful consent. The window will not remain open indefinitely.

CHAPTER 54

Algorithmic Bias in Neurotechnology

The convergence of artificial intelligence with neurotechnology creates a double-edged sword. AI is what makes modern neurotechnology powerful—enabling the decoding of neural signals, the prediction of mental states, and the optimization of brain-computer interfaces. But AI also introduces a profound and often invisible danger: algorithmic bias.

The United Nations report on neurotechnology (A/HRC/57/61) explicitly warns that "issues resulting from the convergence with artificial intelligence, such as algorithmic bias and lack of traceability and transparency, will also be exacerbated through their use." This is not a theoretical concern. It is a structural feature of how AI systems learn and operate.

How Bias Enters the System

Algorithmic bias in neurotechnology does not emerge from malice. It emerges from data. AI systems are trained on datasets that reflect existing societal patterns, including patterns of discrimination, inequality, and exclusion. When those datasets are used to train algorithms that interpret neural signals, the biases embedded in the data become embedded in the technology.

Consider a brain-computer interface designed to detect cognitive engagement or emotional state. If the training data is drawn predominantly from one demographic group—say, young, healthy, male subjects from wealthy countries—the algorithm will learn to recognize neural patterns associated with that group. When the same device is used on a person from a different demographic background—an older woman, a person with a neurological condition, or someone from a different cultural context—the algorithm may misinterpret their neural signals. It may classify normal brain activity as abnormal, or fail to detect genuine cognitive states.

The UN report highlights that "groups experiencing structural discrimination and marginalized people or those in positions of socioeconomic disadvantage are the most at risk of misuse of neurotechnologies." This risk is amplified when the algorithms themselves are biased.

Vulnerable Populations at Greatest Risk

The consequences of algorithmic bias in neurotechnology fall disproportionately on already vulnerable populations. The UN report identifies several groups that face particular risks.

Persons with disabilities are among the most exposed. The report notes that "neurotechnologies are often developed without proper human rights guidelines and monitoring, and this may expose them to a disproportionate risk of physical and/or mental harms." When AI systems are trained on data from able-bodied populations, they may fail to accurately interpret neural signals from individuals with neurological conditions, movement

disorders, or cognitive disabilities. This can lead to misdiagnosis, inappropriate treatment recommendations, or denial of access to assistive technologies.

Children face unique vulnerabilities. Their brains are still developing, and their neural patterns are different from adults. An AI system trained on adult brain data may misinterpret a child's neural activity, leading to incorrect assessments of attention, learning capacity, or emotional state. The UN report emphasizes that "negative impacts on mental health of children, older persons and persons with disabilities should be considered together with the opportunities."

Older adults also present neural patterns that differ from the young, healthy populations typically used in training datasets. Age-related changes in brain activity, cognitive processing, and neural connectivity can be misread by algorithms that have not been trained on diverse age groups.

People from the global South are systematically excluded from neurotechnology development. The UN report notes that "measures have to be put in place to overcome existing structural factors and extend the benefits of neurotechnologies to people living in the global South." When neurotechnology algorithms are trained primarily on data from wealthy countries, they may not accurately interpret neural signals from populations with different genetic backgrounds, dietary patterns, or environmental exposures.

The Feedback Loop of Discrimination

Algorithmic bias in neurotechnology creates a dangerous feedback loop. When biased algorithms are used to make decisions about individuals—determining their fitness for employment, their eligibility for insurance, their cognitive abilities in educational settings—those decisions can reinforce and amplify existing discrimination.

The UN report warns that "neuroprofiles may be used to discriminate among people on the basis of their mental signatures, leading to decisions (i.e. in relation to the workplace or insurance coverage) that may have a negative impact on the interests and rights of the data subject." A biased algorithm that systematically misclassifies neural signals from certain populations can lead to those populations being denied jobs, insurance, or educational opportunities. This creates a self-perpetuating cycle: the people excluded by biased algorithms are less likely to be included in future training datasets, ensuring that the bias persists and deepens.

The Opacity Problem

A critical dimension of algorithmic bias in neurotechnology is the lack of transparency. Modern AI systems, particularly deep learning models, are often "black boxes"—their internal decision-making processes are opaque even to their developers. The UN report identifies "lack of traceability and transparency" as a key concern.

When a neurotechnology device makes a decision based on neural data—classifying a person's mental state, adjusting a stimulation parameter, or generating a recommendation—it is often impossible to determine why that decision was made. Was the algorithm

responding to a genuine neural signal, or was it misinterpreting noise? Was it applying a biased pattern learned from unrepresentative training data? Without transparency, these questions cannot be answered.

This opacity has profound implications for accountability. If a biased algorithm leads to discrimination, who is responsible? The developer who trained the model? The company that deployed it? The clinician or employer who relied on its output? The UN report calls for "enhanced mechanisms" to address these challenges, but the technology is advancing faster than the legal frameworks needed to govern it.

The Need for Inclusive Design

Addressing algorithmic bias in neurotechnology requires a fundamental shift in how these systems are developed. The UN report emphasizes that "without meaningful engagement and participation of persons with disabilities in their design, development and implementation, neurotechnologies risk contributing to stigmatization and objectification, reinforcing structural-based discrimination."

This principle extends beyond disability. Neurotechnology developers must ensure that training datasets are representative of the full diversity of human populations—across age, gender, ethnicity, geography, and neurological status. Algorithms must be tested for bias before deployment, and continuously monitored for discriminatory outcomes after deployment.

The UN report recommends that "human rights impact assessments should also be required" for neurotechnology products. Such assessments must include rigorous analysis of algorithmic bias and its potential effects on vulnerable populations. Without these safeguards, neurotechnology will not be a tool for liberation—it will become a new mechanism for discrimination, hidden behind the opacity of AI and the complexity of the human brain.

CHAPTER 55

AI and Cognitive Liberty — The New Control Challenge

When artificial intelligence makes decisions based on your neural data, a fundamental question arises: who is truly in control? The UN report on neurotechnology and human rights (A/HRC/57/61) identifies this as one of the most pressing challenges of the neurotechnology age. As the report notes, neurotechnologies "generally enable the exposition of cognitive processes" and "bypass the individual's conscious control or awareness." When AI is layered on top of these capabilities, the potential for algorithmic autonomy to override human agency becomes a concrete threat to cognitive liberty.

The Algorithmic Gatekeeper

Modern neurotechnologies do not simply read neural signals and present them to the user. They rely on "artificial intelligence and machine learning, allowing for the processing and analysis of large sets of neurodata." This means that between your brain activity and any action taken based on it, there is an algorithmic intermediary—a system that interprets, filters, and decides.

Consider a brain-computer interface (BCI) designed to control a prosthetic limb. The system collects neural signals, but it is the AI that decodes those signals into intended movements. If the algorithm misinterprets a signal, the limb moves differently than intended. The user experiences a loss of agency—their mental command was not executed as they wished.

Now scale this problem. The UN report describes how neurotechnologies, "with the support of artificial intelligence," can "elaborate sophisticated individualized psychological profiles and use them to predict individuals' behaviour or intentions or to infer their mental characteristics or predispositions." When AI makes inferences about your mental states, and then acts on those inferences without your conscious input, the boundary between your will and the machine's interpretation becomes blurred.

When the Algorithm Decides for You

The most concerning scenarios involve closed-loop systems—neurotechnologies that "read" neural activity and "write" back into the brain without requiring conscious user intervention. The UN report explains that "closed-loop systems result from the fact that both functions, that is reading and writing, are capable of adapting themselves autonomously."

Imagine a neurotechnology device designed to treat depression by detecting patterns of negative thinking and delivering electrical stimulation to interrupt them. If the AI decides that certain thoughts are "negative" and intervenes, who is making the decision about what constitutes acceptable mental content? The user may not even be aware that their thoughts are being monitored and modified.

The report warns that neurotechnologies can "interfere and manipulate individuals" through neuromodulation, altering "the physical and mental processes of a person's inner sphere in ways similar to 'brainwashing.'" When AI drives these interventions autonomously, the manipulation becomes continuous, adaptive, and invisible to the individual.

The Erosion of Informed Consent

A critical dimension of the AI-cognitive liberty challenge is consent. The UN report observes that "neurotechnologies allow the level of conscious reasoning to be bypassed, leaving individuals without psychological protection from having their mind involuntarily read." When AI systems operate autonomously, users may not even know that decisions are being made based on their neural data.

The report further notes that "individuals may not fully understand the implications of using neurotechnologies" and may "unknowingly surrender control over their cognitive processes." This is particularly acute when AI algorithms are opaque—when users cannot understand how their neural data is being interpreted or what decisions are being made.

The UN report emphasizes that "consent requirements to avoid the unauthorized collection and processing of brain data should also be revisited and reinforced." But consent becomes meaningless when the decision-maker is an algorithm operating beyond human awareness or comprehension.

The Accountability Gap

When AI makes autonomous decisions based on neural data, who bears responsibility for harmful outcomes? The UN report identifies this as a fundamental governance challenge. If a closed-loop system delivers inappropriate stimulation, causing psychological harm, the lines of accountability are unclear. The developer of the AI? The company that deployed the device? The clinician who prescribed it? The user who consented to its use?

This accountability gap threatens cognitive liberty because it removes the possibility of meaningful recourse. If you cannot identify who is responsible for an algorithmic decision that affects your mental states, you cannot challenge that decision or seek redress for harm.

Restoring Human Agency

The UN report implicitly calls for a reassertion of human agency in the neurotechnology-AI nexus. This requires several measures: first, that users must retain ultimate control over neurotechnology systems, with the ability to override or disable autonomous functions. Second, that AI algorithms used in neurotechnology must be transparent and explainable—users must be able to understand how decisions are made. Third, that meaningful consent must be obtained for any autonomous operation, with clear disclosure of what the AI is authorized to do.

The report's framework for freedom of thought is instructive. It states that "everyone has the right to develop thoughts autonomously, free from 'impermissible' external influence." When AI makes decisions based on neural data, it constitutes an external influence on thought

processes. The question is whether that influence is permissible—and that determination must be made with full transparency and respect for human agency.

The battle for cognitive liberty in the age of AI is not simply about preventing unauthorized access to the mind. It is about ensuring that when technology interacts with our neural processes, the final decision—the ultimate expression of will—remains with the human being, not the algorithm.

PART X

Society in the Neurotechnology Age

The United Nations Human Rights Council report of August 2024 (A/HRC/57/61) provides the most authoritative framework yet for understanding how neurotechnology will reshape society. The report's core insight is devastating in its clarity: neurotechnologies are "unique and socially disruptive" because they enable the exposition of cognitive processes, allow direct alteration of mental processes, bypass conscious control, enable non-consensual external access to thoughts, and collect neurodata of "a highly sensitive nature."

The Transformation of Daily Life

The report documents that consumer neurotechnology devices are already on the market—wearable headsets, wireless earbuds with biosensing electrodes, and "cosmetically invisible implants." These are being marketed for cognitive performance, relaxation, sleep, and education. Yet the report warns that "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge." Research on potential side-effects, long-term impacts, and their "possible irreversible character is lacking."

This creates a profound societal challenge: millions of people may adopt technologies that alter their neural functioning without understanding the consequences. The report notes that neurotechnology companies are "already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future."

The Erosion of the Inner Fortress

The report frames the human rights implications in stark terms: "For a long time, human mind and conscience were considered the container of the inner self, a fortress, inaccessible from the outside. Penetrating the brain, gaining access to inner mental processes, or directly altering them, represented a frontier not to be crossed."

Neurotechnology changes this. The report identifies that neurotechnologies can serve as "mindreading-like devices" to make inferences about individuals. With AI support, they can "elaborate sophisticated individualized psychological profiles and use them to predict individuals' behaviour or intentions or to infer their mental characteristics or predispositions." This capability threatens to expose what the report calls the *forum internum*—the inner psychological realm where thoughts and convictions are formed.

New Forms of Discrimination

The report warns that "neuroprofiles may be used to discriminate among people on the basis of their mental signatures, leading to decisions (i.e. in relation to the workplace or insurance coverage) that may have a negative impact on the interests and rights of the data subject." This represents a new axis of inequality—a "neuronal divide" where individuals are sorted, judged, and limited based on the patterns of their neural activity.

The Manipulation Economy

Perhaps most disturbingly, the report documents how neurotechnologies enable "behavioural manipulation of individuals by private actors, such as marketing engineers or political campaigners." Neuromarketing strategies can "condition the forming of opinions, as well as influencing an individual's decision-making processes." The report specifically warns that "neurogaming devices can also drive compulsive or addictive use, particularly when neural reward circuits target altering consumption behaviours."

The Vulnerable and the Marginalized

The report identifies that "groups experiencing structural discrimination and marginalized people or those in positions of socioeconomic disadvantage are the most at risk of misuse of neurotechnologies." Persons with disabilities face particular exposure, as neurotechnologies are "often developed without proper human rights guidelines and monitoring." The report insists that persons with disabilities must be involved in the design and development of products that affect them.

A Rapidly Growing Industry in a Deregulated Environment

The report documents explosive industry growth—from $331 million to $7.3 billion in just ten years. Yet "in contrast to the medical development of neurotechnologies, which are generally regulated by established frameworks and processes, the consumer sector remains highly deregulated and is characterized by rapid expansion." The report warns that "without the necessary guardrails, the industry will continue growing unfettered in the same

direction: prioritizing profitability and convenience over ethical and human rights considerations."

This is the societal landscape we face: a technology of unprecedented power, commercialized at breakneck speed, with inadequate safeguards, and with the most vulnerable populations at greatest risk. The UN report's conclusion is that "when dealing with highly disruptive technologies anticipation is key"—but anticipation requires action, not merely awareness.

CHAPTER 56

The Neuronal Divide — A New Class of Cognitive Haves and Have-Nots

The United Nations report on neurotechnology warns of a future already taking shape: a world divided not by wealth, geography, or education alone, but by access to the technology that enhances the very capacity to think. This is the Neuronal Divide — a new axis of inequality that cuts deeper than any that has come before.

The Emerging Two-Tier Society

The UN report documents that neurotechnology investments have surged from $331 million to $7.3 billion in just ten years. The industry is being propelled by public funding from large-scale "brain initiatives" launched by the United States, Australia, China, the European Union, Israel, Japan, and the Republic of Korea since 2013. Yet the report notes that "the growing role of the private sector is changing the dynamics of the industry, which is leaning towards the production of direct-to-consumer devices."

The implications are stark. Those who can afford cognitive enhancement will gain compounding advantages in education, employment, and social mobility. Those who cannot will fall further behind — not merely in economic terms, but in their fundamental capacity to compete in a world where enhanced cognition becomes the baseline for professional success.

The report explicitly identifies this risk: "Access to 'neuroenhancement' is often framed in terms of 'equality' and 'democratization,' a human rights approach requires granting priority access to those in need." But the reality is that market forces, not human rights principles, are driving distribution.

Who Gets Enhanced, Who Gets Left Behind

The report documents that companies are already advertising "cosmetically invisible implants that may allow users to control computers or mobile devices from any location." They claim to be working on the "cognitive enhancement, pleasure or leisure of healthy people." Meanwhile, the report warns that "cutting-edge treatments, such as implants controlled by brain-computer interface software, are often costly and require highly specialized surgeries, which may not always be covered by public health."

This creates a perverse dynamic. The report notes that "paradoxically, once neurotechnologies are safe, effective and secure the members of those groups are usually excluded from or discriminated against in accessing those technologies, which may not be affordable." The most vulnerable — those who could benefit most from assistive neurotechnologies — are the least likely to receive them.

The report further warns that "persons with disabilities, in particular, must be granted access to affordable assistive technologies without discrimination." But it acknowledges that "universal access to advanced technology depends, however, on needs assessments and

budgetary decisions." When budgets are constrained and market forces prevail, cognitive enhancement becomes a luxury good.

The Compounding Advantage

The Neuronal Divide is not static. It compounds over time. The report documents how neurotechnologies can improve "cognitive performance" and aid "concentration, relaxation or sleep." A child with access to neuroeducation tools gains learning advantages that accumulate year after year. An adult with cognitive enhancement can outperform peers in knowledge-intensive industries. The enhanced become more productive, earn more, and can afford better enhancement for their children — creating a hereditary cognitive elite.

The report warns that "neuroprofiles may be used to discriminate among people on the basis of their mental signatures, leading to decisions (i.e. in relation to the workplace or insurance coverage) that may have a negative impact on the interests and rights of the data subject." This discrimination will compound existing inequalities. The unenhanced will not only be less competitive; they will be actively excluded from opportunities based on neural profiling.

The Global Dimension

The Neuronal Divide operates at global scale. The report notes that "at the international level, measures have to be put in place to overcome existing structural factors and extend the benefits of neurotechnologies to people living in the global South." Without such measures, the divide between developed and developing nations will widen dramatically.

The report documents that "nearly half of all worldwide patent applications are registered in the United States." The countries that lead in neurotechnology development will control the means of cognitive enhancement. They will set the standards, own the patents, and determine who gets access. The global South risks becoming a permanent underclass in the cognitive age.

The Human Rights Imperative

The UN report grounds its analysis in the foundational principle of human dignity. It warns that neurotechnologies risk "the objectification and instrumentalization of individuals" and may "hinder individuals' capacity for self-determination." When cognitive enhancement becomes a commodity, human worth becomes tied to cognitive capacity — a direct assault on the principle that all humans possess equal dignity regardless of mental ability.

The report insists that "respect for that overarching principle should be monitored throughout the life cycle of neurotechnologies. A human rights-based approach also requires limiting the implementation of applications that are unable to comply with it." This means that the Neuronal Divide is not merely an ethical concern — it is a human rights violation in waiting.

The question is not whether the Neuronal Divide will emerge. The technology is already being commercialized. The investments are already flowing. The question is whether we will act to ensure that cognitive enhancement serves human dignity rather than entrenching

inequality. The UN report's warning is clear: without deliberate intervention, we are building a world where the right to think — to be fully human — becomes a privilege of the few.

CHAPTER 57

What Rights Are We Willing to Give Up?

The UN report on neurotechnology and human rights presents a stark question that every society must now confront: which of our fundamental rights are we prepared to sacrifice in exchange for convenience, safety, productivity, or profit? The report documents that the neurotechnology industry has grown from $331 million to $7.3 billion in just ten years, driven by "the growing role of the private sector" and a trend "towards the production of direct-to-consumer devices." This commercial pressure is already forcing trade-offs that most citizens have not consciously chosen.

Privacy for Convenience

The most immediate trade-off involves mental privacy. The UN report notes that "non-invasive neurotechnologies are used for a variety of medical purposes, but the same devices are also being commercialized for different applications in the area of well-being, in the workplace or as an educational tool." Consumers are already exchanging access to their neural data for the ability to track concentration levels, improve sleep, or enhance cognitive performance. The report warns that "users can wirelessly track their levels of concentration and check their results on their phone" — but this convenience comes at a cost.

The report documents that "consumer neurotechnology companies are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future." When users agree to terms of service they have not read, they surrender control over data that reveals "sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions." The trade-off appears voluntary, but the report notes that "individuals may not fully understand the implications of using neurotechnologies" and "may unknowingly surrender control over their cognitive processes."

Autonomy for Safety

The second trade-off involves personal autonomy in exchange for security. The UN report identifies that "some proposed applications of neurotechnologies in the area of justice and national security may lead to the criminalization of mere thought." With neurotechnology, "people can be punished for unexpressed thoughts or intentions and pre-emptive decisions affecting their lives or interest can be taken even before they translate into action."

This represents a fundamental shift in how societies balance liberty against security. The report specifically warns that "the trend to introduce national security and law enforcement exceptions in artificial intelligence instruments and regulations raises concerns as they may not offer sufficient protection against problematic uses of neurotechnologies." Citizens may accept brain-reading in airports or at border crossings, believing it protects them from threats. But the report notes that such technologies "can serve as 'mindreading-like' devices

to make inferences about individuals" and "elaborate sophisticated individualized psychological profiles" that can be used to "predict individuals' behaviour or intentions."

Cognitive Liberty for Productivity

The third trade-off — perhaps the most insidious — involves surrendering cognitive liberty for workplace productivity. The UN report warns that neurotechnologies "can be used to interfere and manipulate individuals" through devices that "alter the physical and mental processes of a person's inner sphere in ways similar to 'brainwashing.'" In workplace settings, employees may accept neural monitoring as a condition of employment, allowing employers to track attention levels, emotional states, and cognitive performance.

The report documents that "neuromarketing strategies have already demonstrated" how neurotechnology can be "successfully used to condition the forming of opinions, as well as influencing an individual's decision-making processes." When applied in the workplace, this enables "behavioural manipulation of individuals by private actors" on an unprecedented scale. The report notes that "consumer neurotechnology companies" are developing devices that can be used "even on a permanent basis, while performing different activities." The trade-off appears reasonable — increased productivity for a paycheck — but the report warns that "individuals may not fully understand the implications" and that "the level of conscious reasoning can be bypassed, leaving individuals without psychological protection."

Consent Under Duress

The UN report emphasizes that meaningful consent becomes impossible when power asymmetries exist. It notes that "in contexts characterized by power asymmetries, for example in the military domain, human rights concerns may arise with regard to interventions for enhancement purposes." The report specifically recommends that "coercive uses of neurotechnologies to augment resilience on the battlefield (i.e. to reduce emotions, empathy or the need for sleep) should be prohibited."

But the same logic applies in civilian contexts. The report observes that "particularly in the case of consumer-oriented neurotechnology devices, it is hard to see how users can provide informed consent when their potential long-term effects are largely unknown." When individuals cannot understand what they are agreeing to — when the risks are not yet documented, when the long-term impacts on mental health are not studied — consent becomes a legal fiction. The report notes that "research on potential side-effects, long-term impacts and their possible irreversible character is lacking."

The Slippery Slope

The UN report frames these trade-offs not as hypotheticals but as ongoing processes. It warns that "the risk is that, without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations." Each small concession — allowing a device to monitor attention, accepting neural data collection for a discount on insurance, permitting

workplace brain surveillance for productivity bonuses — normalizes the next, larger intrusion.

The report's framework makes clear that freedom of thought is an absolute right, one that "does not allow for any interference whatsoever." Yet the commercial pressures documented in the report are already eroding this protection. The question is not whether trade-offs will be made, but whether they will be made consciously, with full democratic deliberation — or whether they will simply happen, device by device, transaction by transaction, until the fortress of the inner self has been surrendered for the convenience of the moment.

CHAPTER 58

Neurotechnology and Democracy — The Threat to Free Elections

Democracy depends on the premise that citizens form their political opinions freely, deliberate independently, and cast votes based on their authentic beliefs. Neurotechnology now threatens each of these pillars. The UN Human Rights Council Advisory Committee's 2024 report (A/HRC/57/61) identifies the mechanisms by which neurotechnologies "can be used in ways that may erode democracy and the rule of law."

Neuromarketing and Political Manipulation

The report explicitly warns that "as 'neuromarketing' strategies have already demonstrated, [neurotechnologies] can be successfully used to condition the forming of opinions, as well as influencing an individual's decision-making processes." This is not speculative. The report notes that these strategies "based on the way in which individuals respond to certain stimuli, can tailor their messages specifically to certain neurotypes."

In an electoral context, this means political campaigns could use consumer neurotechnology devices — already on the market for "well-being, in the workplace or as an educational tool" — to identify which neural responses particular messages trigger, then micro-target voters with content designed to bypass rational deliberation and exploit emotional vulnerabilities. The report warns that this "enables, to an unprecedented extent, behavioral manipulation of individuals by private actors, such as marketing engineers or political campaigners."

Cognitive Warfare Against Electorates

The convergence of neurotechnology with AI creates even more potent threats. The report observes that "with the support of artificial intelligence and machine learning, allowing for the processing and analysis of large sets of neurodata, the opportunities and uses of neurotechnologies are multiplying." These same tools can be weaponized against democratic processes.

The report identifies a critical danger: "neurotechnologies can be used to interfere and manipulate individuals. Through neuromodulation devices, the physical and mental processes of a person's inner sphere can be altered in ways similar to 'brainwashing.'" Applied at scale during an election cycle, such techniques could alter the political preferences of entire populations without their awareness.

The report specifically warns that "the capabilities offered by neurotechnologies may serve other non-legitimate purposes" beyond commercial neuromarketing. Foreign adversaries could deploy these technologies to destabilize democratic elections, using directed neuromodulation to induce specific emotional states — fear, anger, apathy — that favor particular outcomes.

Neural Surveillance and Voter Intimidation

The report documents how neurotechnologies "enable the exposition of cognitive processes" and "enable non-consensual external access to thoughts, emotions and mental states." In an electoral context, this creates unprecedented possibilities for voter surveillance and intimidation.

The report notes that "neurodata provide highly sensitive information and allow aspects of individuals' identity and personality to be revealed, such as sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions." If voter neural data — collected through consumer devices — could be used to infer political preferences, the secret ballot itself could be undermined. Voters could be coerced into supporting particular candidates, or punished for their actual choices after the fact.

The report warns that "neuroprofiles may be used to discriminate among people on the basis of their mental signatures, leading to decisions (i.e. in relation to the workplace or insurance coverage) that may have a negative impact on the interests and rights of the data subject." In a democracy, the extension of this logic to political participation is obvious: citizens could face discrimination or retaliation based on neural profiles that reveal their political leanings.

The Manipulation of the Forum Internum

At the deepest level, the threat to democracy is a threat to the *forum internum* — the inner realm where thoughts and convictions are formed. The report emphasizes that "freedom of thought (International Covenant on Civil and Political Rights, art. 18) comes to the fore" when neurotechnologies "have the ability to interfere with the human mind."

The report clarifies that "everyone has the right to develop thoughts autonomously, free from 'impermissible' external influence." But neurotechnology makes impermissible influence possible at a scale and precision never before achievable. The report notes that "freedom of thought is a far-reaching and profound right that encompasses all matters" and is "at the origin of the thinking and is the source of decisional processes, engagement and the development of an individual's consciousness."

Because of "the essential role it is called upon to play in democracy," the report states that "freedom of thought does not allow for any interference whatsoever and must be protected unconditionally, which explains why it is characterized as an absolute right."

The Threat to Free Elections

The convergence of these capabilities — neuromarketing, cognitive warfare, neural surveillance, and manipulation of the *forum internum* — creates a systemic threat to free and fair elections. The report warns that "some proposed applications of neurotechnologies in the area of justice and national security may lead to the criminalization of mere thought." In the electoral domain, the equivalent is the criminalization of unexpressed political dissent — or its detection and suppression before it can be expressed at the ballot box.

The report observes that "with the support of neurotechnologies, people can be punished for unexpressed thoughts or intentions and pre-emptive decisions affecting their lives or interest can be taken even before they translate into action." In an election, this could mean pre-emptively disqualifying voters, gerrymandering based on neural profiles, or suppressing turnout among populations whose neural data indicates opposition to the ruling party.

The report's conclusion is stark: "Facing the implications of neurotechnologies may require not only developing specific policies, legislation and instruments but also reinforcing the institutional structure to ensure that prevention and mitigation policies are effectively implemented." For democracy to survive the neurotechnology age, the absolute protection of the *forum internum* must be non-negotiable — and the use of neurotechnology in political campaigns and elections must be subject to the most stringent prohibitions.

CHAPTER 59

Children and Neurotechnology — A Vulnerable Generation

The United Nations report on neurotechnology and human rights (A/HRC/57/61) identifies children as a group requiring special attention, noting that "negative impacts on mental health of children, older persons and persons with disabilities should be considered together with the opportunities." This recognition is not incidental — it reflects a fundamental truth about neurotechnology's interaction with the developing brain.

The Plasticity Paradox

Children's brains are not miniature adult brains. They are organs in a state of rapid development, characterized by extraordinary neuroplasticity — the ability to reorganize neural pathways in response to experience. This plasticity is what allows children to learn languages, develop social skills, and form their identities. But it is also what makes them uniquely vulnerable to neurotechnology.

The report emphasizes that neurotechnologies "can affect mental health and provoke alterations in an individual's personality, psychological balance or sense of self-identity." For a child whose sense of self is still forming, these alterations are not temporary disruptions — they become embedded in the architecture of the developing brain. A neural intervention that might be reversible in an adult can permanently shape a child's cognitive and emotional development.

Neurogaming and the Developing Reward System

The report specifically warns about "neurogaming" devices that "can drive compulsive or addictive use, particularly when neural reward circuits target altering consumption behaviours." For children, whose prefrontal cortex — the brain region responsible for impulse control and decision-making — is not fully developed until their mid-twenties, this vulnerability is magnified.

Consumer neurotechnology products marketed to children and adolescents, such as EEG headsets for "brain training" or "focus improvement," are already entering the market. The report notes that "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge. Moreover, research on potential side-effects, long-term impacts and their possible irreversible character is lacking."

This is a dangerous combination: products that directly interface with the developing brain, marketed with unproven claims, and deployed without adequate safety testing. The report's warning applies with particular force to children: "In the absence of a medical need and without being supported by extensive trials and scientific evidence, devices may become harmful for mental health and their extensive use may lead to public health problems."

Identity Formation Under Surveillance

The report identifies that neurotechnologies "enable the exposition of cognitive processes" and "enable non-consensual external access to thoughts, emotions and mental states." For children, who are still learning what their own thoughts and feelings mean, the presence of external monitoring can fundamentally alter the process of identity formation.

The *forum internum* — the inner realm of thought — is not merely a private space for adults. For children, it is a workshop where identity is constructed, where beliefs are tested, and where the self is discovered. The report's affirmation that "freedom of thought does not allow for any interference whatsoever and must be protected unconditionally" applies to children as fully as to adults.

Yet children are particularly susceptible to the forms of manipulation the report describes. The report warns that neurotechnologies "can be successfully used to condition the forming of opinions, as well as influencing an individual's decision-making processes." When these technologies are deployed in educational settings — to monitor attention, assess learning, or "personalize" instruction — they risk conditioning not just what children learn, but how they think.

The Consent Problem

The report repeatedly emphasizes the importance of "meaningful and effective free and informed consent." But children present a fundamental challenge to this principle. As the report notes, "particularly in the case of consumer-oriented neurotechnology devices, it is hard to see how users can provide informed consent when their potential long-term effects are largely unknown."

For children, the consent problem is compounded. Parents may consent on behalf of their children, but they cannot know the long-term effects of neurotechnology on a developing brain. The report warns that "individuals may not fully understand the implications of using neurotechnologies" and "may unknowingly surrender control over their cognitive processes." Children, by definition, lack the cognitive maturity to understand these implications.

The report's framework for vulnerable groups applies with special force: "Consent to participate in trials may be provided under coercion or without being fully informed or made aware of negative side-effects on, for example, health, privacy and autonomy." When the trial is a classroom activity or a marketed "educational tool," the line between voluntary participation and coercion becomes dangerously blurred.

The Right to an Open Future

The report's analysis of neurotechnology's impact on "personal integrity" raises a profound question for children. It notes that neurointerventions can produce "alterations in an individual's personality, psychological balance or sense of self-identity." For a child, these alterations shape not just who they are, but who they can become.

The report's framework for freedom of thought provides the foundation: "Everyone has the right to develop thoughts autonomously, free from 'impermissible' external influence." For

children, this right to autonomous development is not merely about protecting existing thoughts — it is about preserving the capacity to form thoughts that have not yet been conceived.

This is the principle of the "open future" — the idea that children have a right to develop their own identities, beliefs, and values without being pre-committed by decisions made before they could participate in them. Neurotechnology, by shaping the developing brain, can close off futures before a child has the chance to choose them.

Special Protections Required

The report calls for "context-specific standards to protect against the non-consensual intrusion of third parties" and emphasizes that "consent requirements to avoid the unauthorized collection and processing of brain data should also be revisited and reinforced." For children, these standards must be even more stringent.

The report's framework for "groups in vulnerable situations" provides the model. It notes that "persons with disabilities are particularly exposed to medical and scientific procedures, research and experimentation" and that "neurotechnologies are often developed without proper human rights guidelines and monitoring." The same applies to children, who are exposed to neurotechnology in educational and commercial settings without adequate protections.

The report's conclusion applies with particular urgency to children: "States have a duty to ensure that neurotechnology products, for both medical and consumer use, are developed following strict safety and security controls." For children, these controls must include explicit prohibitions on non-medical use, mandatory long-term safety studies, and a presumption against any neurotechnology that could alter the developing brain.

The generation growing up with neurotechnology will not have the chance to choose whether they want their brains interfaced with devices, their thoughts monitored by algorithms, or their identities shaped by neural interventions. The choice must be made now — and it must be made on their behalf, with their future as the only consideration.

CHAPTER 60

Workplace Neurosurveillance — The Coming Cognitive Panopticon

The employment relationship has always involved some surveillance of worker activity. Time clocks, productivity metrics, security cameras, and electronic communication monitoring have been the established tools. Neurosurveillance — the direct monitoring of brain activity in the workplace — represents a qualitative leap. It moves the surveillance frontier from behaviour to cognition, from what workers do to what they think.

The Technology Is Already Here

Neurosurveillance in the workplace is not speculative. Commercial products are already deployed. EEG headbands marketed for productivity monitoring track attention levels in real time and report them to employers. SmartCap, deployed in mining, transport, and aviation, monitors fatigue through EEG and reports drowsiness to supervisors. Emotiv's enterprise products provide aggregate workforce attention metrics to managers. Chinese companies are reported to have deployed brain-monitoring caps to factory workers, with the data used to optimize productivity and detect emotional states.

The technologies are non-invasive, increasingly inexpensive, and integrated into ordinary workplace equipment such as helmets and hats. Their adoption is being driven by claims of safety enhancement, productivity gain, and reduced absenteeism. The marketing emphasises benevolent applications; the rights implications are largely absent from the corporate discourse.

Documented Deployments

SmartCap (Australia/global): EEG-based fatigue monitoring deployed in mining, road transport, and aviation. Used by Rio Tinto, BHP, and major transport firms.

Hangzhou Zhongheng Electric (China): Reported deployment of brain-monitoring caps to assembly-line workers, with data feeding into productivity systems.

Emotiv Insight Enterprise: Attention and engagement monitoring marketed to corporate clients.

Neuroprofiler / Neurolinx: Neural marketing analytics deployed in retail and financial services.

Various pilot programmes: Reported in financial trading, professional sports, and military training contexts.

The Inevitability Argument

Proponents of workplace neurosurveillance argue that the technology is inevitable and that responsible deployment is preferable to prohibition. The argument follows a familiar pattern: the technology exists, competitors will adopt it, the only choice is between thoughtful and reckless deployment. Workers benefit from improved safety. Companies benefit from improved productivity. Society benefits from reduced workplace accidents.

The argument is structurally similar to the arguments deployed in prior expansions of workplace surveillance — each of which has, in retrospect, produced harms that the proponents did not anticipate or chose not to discuss. The cumulative effect of decades of incremental surveillance expansion is a workplace environment in which worker autonomy is substantially diminished from historical norms. Neurosurveillance is not an isolated step; it is the logical continuation of a multi-decade trajectory.

Why Consent Cannot Be Free

The standard legal protection for personal data is consent. Workers can in principle consent to neurosurveillance, and the technology can be lawfully deployed with that consent. The structural problem is that consent in the employment relationship is rarely free.

A worker presented with a neurosurveillance requirement at the time of hiring, or as a condition of continued employment, faces an asymmetric choice: accept the surveillance and keep the job, or decline and accept unemployment. The European Data Protection Board has recognised this asymmetry, noting that consent in the employment context is rarely freely given. The same reasoning applies with particular force to neural data, the most intimate category of personal information.

Several jurisdictions have begun to address this. The proposed EU AI Act prohibits the use of AI systems to infer emotions in the workplace and educational contexts — a partial restriction that addresses some but not all neurosurveillance applications. The Illinois Biometric Information Privacy Act provides stronger protections, though it does not specifically address neural data. Most jurisdictions remain entirely unprotected.

The Discrimination Dimension

Workplace neurosurveillance creates new vectors of discrimination. Neural data can reveal characteristics that workers have legitimate reasons to keep private: neurological conditions, mental health status, cognitive style, and potentially pregnancy and substance use. An employer with access to neural data may discriminate — explicitly or implicitly — against workers whose neural signatures suggest disfavoured characteristics.

The discrimination need not be conscious. AI-driven HR systems trained on neural data may produce systematically disparate outcomes for protected groups, without any individual decision-maker forming a discriminatory intent. The discrimination becomes invisible, structural, and resistant to traditional legal remedies.

The Chilling Effect on Cognition

A deeper concern is the effect of pervasive cognitive monitoring on the workers themselves. Workers aware that their attention is monitored, their fatigue tracked, and their emotional states logged may alter their behaviour to optimise the measured metrics. They may suppress emotional responses, force attention even when fatigue is genuine, and avoid mental activities (such as creative wandering) that produce neural signatures interpretable as unproductive.

The result is a form of cognitive performativity that may be subtly destructive to mental health and creativity. The worker is not just observed but altered by the observation — the classical panopticon effect, applied not to physical behaviour but to mental life. The cumulative effect of decades of working under cognitive surveillance has no precedent and cannot be fully predicted.

The Required Response

A response adequate to workplace neurosurveillance requires several elements:

33. **Prohibition of cognitive monitoring as a condition of employment.** Workers should not be required to accept neural surveillance as a condition of hiring, retention, or advancement. The asymmetric nature of the employment relationship makes consent insufficient protection.
34. **Strict purpose limitation.** Where neural data is collected with genuine consent for safety purposes (such as detecting fatigue in hazardous occupations), use of that data for any other purpose should be prohibited.
35. **Worker representation in deployment decisions.** Trade unions, works councils, and equivalent worker representative bodies should have substantive rights to negotiate over the deployment of neurosurveillance technologies.
36. **Independent auditing.** Where neurosurveillance is deployed, independent technical auditing should verify that the system functions as represented and that data flows comply with applicable law.
37. **Personal access and deletion rights.** Workers should have ongoing rights to access neural data collected about them, to obtain explanations of inferences drawn, and to require deletion of data upon termination of employment.

The Stakes

The workplace is the setting in which most adults spend most of their waking hours. The conditions of work shape the conditions of life. The introduction of cognitive surveillance into the workplace, without adequate protection, would constitute one of the most significant expansions of monitoring in modern history — with consequences that would extend far beyond the workplace itself, into the broader social texture and the development of human autonomy.

The decisions made in the next several years, by legislators, regulators, and employers, will determine whether the workplace becomes a cognitive panopticon or a setting in which the protections of neuro-rights are operationalised in concrete and enforceable terms.

PART XI

Who Benefits Commercially? — The Neurotechnology Industry

The United Nations report is unequivocal: neurotechnology is "a rapidly growing industry in a highly deregulated environment." The numbers tell the story of an explosion. Private investment in neurotechnology has surged from $331 million to $7.3 billion in just ten years. Large-scale "brain initiatives" launched by the United States, Australia, China, the European Union, Israel, Japan, and the Republic of Korea since 2013 have poured public funding into the sector. The result is an industry that is "leaning towards the production of direct-to-consumer devices" — devices that bypass the regulatory safeguards that govern medical neurotechnology.

The Private Sector Takes the Lead

The report notes that "the growing role of the private sector is changing the dynamics of the industry." With "breakthrough advances reported daily," the pressure from large technology companies to place neurotechnology products on the market "will inevitably increase." Nearly half of all worldwide patent applications in neurotechnology are registered in the United States alone. Companies are already advertising "cosmetically invisible implants" that allow users to control computers from any location, and they "claim to be working on the cognitive enhancement, pleasure or leisure of healthy people."

A Regulatory Vacuum

The contrast between medical and consumer neurotechnology could not be starker. Medical neurotechnologies are "generally regulated by established frameworks and processes." The consumer sector, by contrast, "remains highly deregulated and is characterized by rapid expansion." The report warns that "the possibility that, in the coming years, those products with inadequate safety measures and unclear or underestimated human rights risks may be widely commercialized is real."

The consequences are predictable. "Existing loopholes in regulations, lack of technical expertise and capacity and the absence of adequate oversight bodies are factors that will certainly be exploited by large companies seeking profits." The risk is that "without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations."

The Data Gold Rush

Neurodata — defined as "information gathered from the brain and/or from the nervous system, including inferences extracted from it" — is the industry's raw material. The report observes that "due to its biometric character, commercial interest in data collected through medical and non-medical neurotechnology devices will grow exponentially." Consumer neurotechnology companies are "already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future."

The report notes that "data can be sold to third parties without the knowledge or authorization of users." Users may be "encouraged to exchange their data for the use of services on an (allegedly) voluntary basis." The asymmetry is profound: individuals may not fully understand what they are surrendering when they put on a headband or implant a chip.

The Market's Direction

The report identifies a clear trajectory. "With the support of artificial intelligence and machine learning, allowing for the processing and analysis of large sets of neurodata, the opportunities and uses of neurotechnologies are multiplying." Non-invasive biosensors integrated into "easy-to-wear headsets" are already being promoted for applications that may be used "even on a permanent basis, while performing different activities." The sensor systems of the new generation of wireless earbuds "include biosensing electrodes, enabling the measurement of various biosignals, including brain activity."

Yet the report notes a troubling gap: "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge." Research on "potential side-effects, long-term impacts and their possible irreversible character is lacking."

The Military-Industrial Connection

The report does not explicitly name the military-industrial complex, but the implications are clear. The same technologies being developed for consumer markets — neural interfaces, brain-computer connections, cognitive enhancement — have parallel applications in national security. The report warns of "the trend to introduce national security and law enforcement exceptions in artificial intelligence instruments and regulations," noting that such exceptions "raise concerns as they may not offer sufficient protection against problematic uses of neurotechnologies."

Who Benefits?

The answer, according to the report, is clear: the companies that control the data and the devices. The $7.3 billion industry is driven by private investment, patent filings, and the race to market. The benefits flow to those who can commercialize neurotechnology fastest, with the fewest regulatory constraints, and with the greatest access to the data that powers their systems. The costs — in privacy, autonomy, and human rights — are externalized to users, to vulnerable populations, and to future generations.

CHAPTER 61

The Neurotechnology Market — $7.3 Billion and Growing

The numbers are staggering, and they are accelerating. According to the United Nations Human Rights Council Advisory Committee's 2024 report, private investments in neurotechnology have increased from $331 million to $7.3 billion in just ten years. This is not a niche industry. This is a market that has grown by more than twenty-fold in a single decade, and it shows no signs of slowing.

The Investment Surge

The report attributes this explosive growth to two primary drivers: large-scale "brain initiatives" launched by multiple countries since 2013, and the increasing role of the private sector. After the United States initiated its BRAIN Initiative, Australia, China, the European Union, Israel, Japan, and the Republic of Korea developed similar research projects. These government-funded programs created the foundational science. Private capital then rushed in to commercialize it.

The result is a market that is "leaning towards the production of direct-to-consumer devices." The report notes that "with breakthrough advances reported daily, the pressure from large technology companies to place neurotechnology products in the market will inevitably increase." Nearly half of all worldwide patent applications in neurotechnology are registered in the United States alone.

The Deregulated Consumer Sector

The most striking finding of the UN report is the asymmetry between the medical and consumer sectors. Medical neurotechnologies are "generally regulated by established frameworks and processes." Clinical trials, safety protocols, and ethical oversight govern their development. But the consumer sector is characterized by "rapid expansion" in a "highly deregulated environment."

This is where the money is flowing. Non-invasive neurotechnologies—wearable headsets, earbuds with biosensing electrodes, wristbands, electrode patches, and even temporary tattoos—are being commercialized for "well-being, in the workplace or as an educational tool." The most popular devices promise to improve "cognitive and sporting performance" and aid "concentration, relaxation or sleep." Users can "wirelessly track their levels of concentration and check their results on their phone."

But the report issues a stark warning: "the effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge." Research on "potential side-effects, long-term impacts and their possible irreversible character is lacking."

The Data Economy

The true value of the neurotechnology market may not be in the devices themselves, but in the data they generate. The UN report explains that neurotechnologies are "nurtured by 'neurodata,' which are needed for their own functioning, calibration and optimization." These data are defined as "the information gathered from the brain and/or from the nervous system, including inferences extracted from it."

The commercial interest in this data "will grow exponentially." Consumer neurotechnology companies are "already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future." The report warns of "predatory practices" and "hacking" as significant risks.

The Patent Race

The UN report notes that nearly half of all worldwide patent applications in neurotechnology are registered in the United States. This is not merely a statistic; it is a strategic indicator. Patent filings reveal where companies believe the future lies. They are betting on brain-computer interfaces, neural implants, and the algorithms that decode neural signals.

The report describes a company that "has recently developed a secure interface for communication 'with the power of thought'" and is "conducting large-scale trials of this technology, which can be implanted into the brain through blood vessels." Other companies advertise "cosmetically" invisible implants that "may allow users to control computers or mobile devices from any location." They claim to be working on "the cognitive enhancement, pleasure or leisure of healthy people."

The Regulatory Vacuum

The report paints a concerning picture of what happens when a $7.3 billion market operates without adequate guardrails. "Existing loopholes in regulations, lack of technical expertise and capacity and the absence of adequate oversight bodies are factors that will certainly be exploited by large companies seeking profits."

The risk is not hypothetical. The report states plainly: "the possibility that, in the coming years, those products with inadequate safety measures and unclear or underestimated human rights risks may be widely commercialized is real." They "may become pervasive throughout daily life despite the fact that, in most countries, applicable regulations are unclear, weak or non-existent."

The conclusion is unavoidable: "Without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations."

The Market's Blind Spot

The UN report identifies a fundamental tension at the heart of the neurotechnology market. The industry's growth depends on access to neural data. But that same access creates

profound human rights risks. The report identifies “freedom of thought,” “right to privacy,” “right to personal integrity,” “right to health,” and “prohibition of torture” as rights particularly at risk.

The market, left to itself, will not address these risks. The report calls for “context-specific standards” and “enhanced protective frameworks.” But the market is moving faster than the law. The $7.3 billion industry is growing while regulators are still debating whether neural data deserves special protection.

The question is not whether the market will continue to grow. It will. The question is whether the growth will be guided by human rights or by profit alone. The UN report makes clear that the answer, as of 2024, is leaning dangerously toward the latter.

CHAPTER 62

Big Tech's Brain Rush — Neuralink, Synchron, and the Silicon Valley Players

The UN report's warning about a $7.3 billion market growing "unfettered" is not an abstraction. It is a description of a race already underway—a race being run primarily by a handful of well-funded companies in Silicon Valley and beyond. These are the players that will shape the commercial neurotechnology landscape, for better or worse.

Neuralink: The Most Visible, The Most Controversial

Elon Musk's Neuralink has captured more public attention than any other neurotechnology company. The UN report notes that Neuralink "claims to have developed a brain implant that will allow a connection to the Internet at the biological level, which will allow people to communicate wirelessly with anyone with a similar implant and computer setup."

The company has already begun human trials. The UN report confirms that Neuralink "has recently started trials of experimental implantable chips in humans." Its stated ambitions extend beyond medical applications. The company "already advertise[s] 'cosmetically' invisible implants that may allow users to control computers or mobile devices from any location." They claim to be working on "the cognitive enhancement, pleasure or leisure of healthy people."

Neuralink's business model is classic Silicon Valley: move fast, break things, and figure out regulation later. The company's invasive approach—requiring surgical implantation of electrodes directly into brain tissue—offers higher precision and power than non-invasive alternatives. But it also entails greater risks. The UN report notes that "invasive applications and procedures are more precise and powerful but entail more risks as they require the surgical implantation of electrodes or implants."

Synchron: The Stealth Challenger

While Neuralink dominates headlines, Synchron has been quietly advancing a less invasive approach. The UN report notes that "a company that has recently developed a secure interface for communication 'with the power of thought' is conducting large-scale trials of this technology, which can be implanted into the brain through blood vessels."

That company is Synchron. Its stentrode device is delivered to the brain through the jugular vein, avoiding open brain surgery. This approach, while less precise than Neuralink's direct cortical implants, dramatically reduces surgical risk. Synchron's device is already in human trials, and the company has received FDA approval for its investigational device exemption.

The competitive dynamic between Neuralink and Synchron illustrates a broader tension in the industry: the trade-off between precision and safety. More invasive technologies offer greater capability but carry higher risks. Less invasive approaches are safer but may be less

powerful. The UN report's framework for distinguishing between invasive and non-invasive technologies provides a useful lens for understanding these trade-offs.

Kernel: The Non-Invasive Alternative

Not all major players are pursuing implants. Kernel, a Los Angeles-based company, has focused on non-invasive neuroimaging devices. Its Kernel Flow headset uses functional near-infrared spectroscopy (fNIRS) to measure brain activity through the skull, without requiring surgery.

Kernel represents a different business model: direct-to-consumer neurotechnology. The UN report notes that "non-invasive neurotechnologies are used for a variety of medical purposes, but the same devices are also being commercialized for different applications in the area of well-being, in the workplace or as an educational tool."

The report is skeptical of such claims. "The effectiveness of most of those commercialized products is neither proved nor backed up by scientific evidence/knowledge," it states. "Moreover, research on potential side-effects, long-term impacts and their possible irreversible character is lacking."

Kernel's approach raises the same questions that apply to the entire consumer neurotechnology sector: Are these devices actually doing what they claim? And what happens to the neural data they collect?

NeuroSky: The Pioneer

NeuroSky, founded in 2004, was one of the first companies to bring consumer EEG devices to market. Its headsets, which use a single dry electrode placed on the forehead, have been used in education, gaming, and meditation applications.

NeuroSky's business model has been built on licensing its technology to other companies rather than selling directly to consumers. Its chips have been embedded in toys, educational tools, and wellness devices. The company represents the less glamorous but more pervasive end of the neurotechnology market: not revolutionary brain implants, but modest wearables that collect and analyze neural data.

The UN report's warning about data collection applies directly here. "Consumer neurotechnology companies are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future."

The Common Thread: Data

Despite their different approaches, these companies share a fundamental characteristic: their business models depend on access to neural data. The UN report defines "neurodata" as "the information gathered from the brain and/or from the nervous system, including inferences extracted from it."

The report identifies a clear risk: “Neurodata provide highly sensitive information and allow aspects of individuals’ identity and personality to be revealed, such as sexual orientation, personality traits and cognitive performance, and mental states, such as intentions, beliefs and emotions.”

The companies themselves may not intend to abuse this data. But the report warns that “data can be sold to third parties without the knowledge or authorization of users.” And “neuroprofiles may be used to discriminate among people on the basis of their mental signatures.”

The Ethical Vacuum

The UN report’s assessment of the industry is stark. “The growing role of the private sector is changing the dynamics of the industry, which is leaning towards the production of direct-to-consumer devices.” And “the pressure from large technology companies to place neurotechnology products in the market will inevitably increase.”

The report identifies a critical gap: “in contrast to the medical development of neurotechnologies, which are generally regulated by established frameworks and processes, the consumer sector remains highly deregulated and is characterized by rapid expansion.”

Neuralink, Synchron, Kernel, and NeuroSky are not villains. They are companies operating in a regulatory vacuum, pursuing legitimate business opportunities. But the UN report makes clear that this vacuum is dangerous. Without guardrails, the industry will continue to prioritize “profitability and convenience over ethical and human rights considerations.”

The question is not whether these companies will succeed. They already are. The question is whether the success will come at the cost of the rights the UN report identifies as “particularly at risk”: freedom of thought, privacy, personal integrity, and the prohibition of torture. The answer, as of 2024, is still being written—by the companies themselves, and by the regulators who have yet to catch up.

CHAPTER 62-bis

The Device Makers Atlas — A Complete Industry Overview

Chapter 62 introduced the most visible Silicon Valley players. This chapter takes a more comprehensive view: it surveys the global device-maker landscape across three tiers — implantable/surgical, clinical-grade non-invasive, and consumer wearable — and asks of each major company what their devices actually do, what they can read or write, what data they retain, and what their public roadmap is. The atlas is necessarily a snapshot; the field is moving fast. But a legislator drafting today needs to know which companies' products will be on shelves and in operating theatres tomorrow, and what those products will be capable of. We have deliberately included companies of widely different scale, from billion-dollar implant developers to two-person wearable start-ups, because the legislative challenge is exactly that range: a single regulatory regime must apply to both.

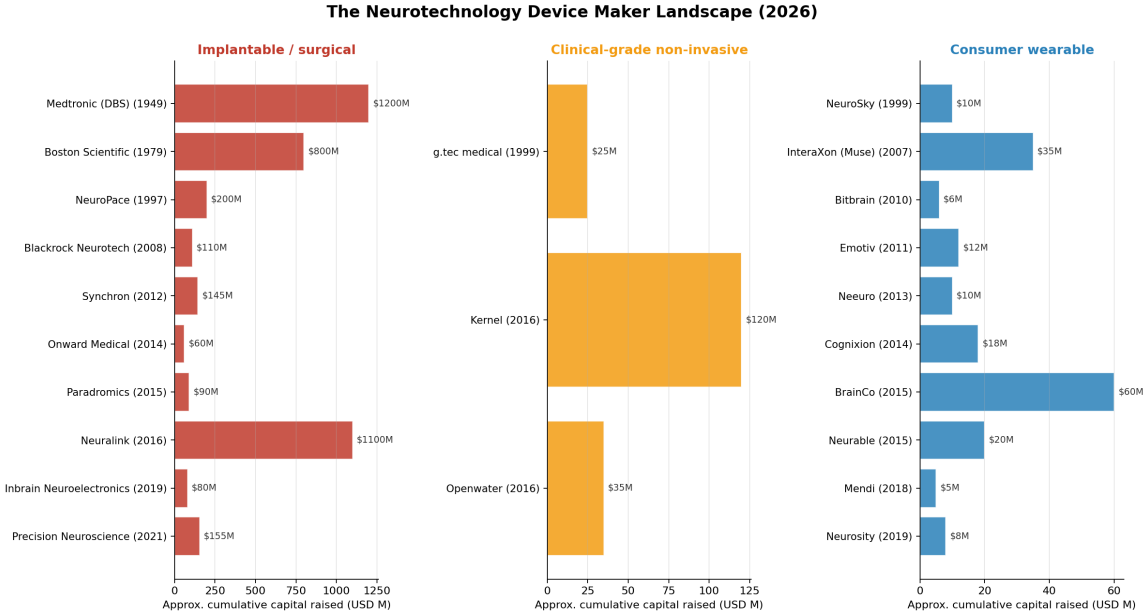


Figure 62bis.1. The neurotechnology device-maker landscape in 2026, organised by invasiveness tier and approximate cumulative capital raised. Note that the largest tier by capital is the implantable category, but the largest tier by population reached is the consumer wearable category. Privacy risk does not track surgical risk: a meditation headband in a school can affect more minds than a cortical implant in a hospital.

Tier 1 — Implantable and Surgical Neurotechnology

Medtronic

Medtronic, founded in 1949 and headquartered in Minneapolis, is the world's largest medical-device company by revenue and the dominant supplier of deep-brain stimulation systems. Its Activa and Percept families of DBS implants are FDA-approved for Parkinson's disease, essential tremor, dystonia and treatment-resistant epilepsy, with humanitarian device exemption use in obsessive-compulsive disorder. The Percept PC introduced sensing (closed-loop) capability in 2020; the Percept RC, released in 2024, extended battery life and adaptive stimulation. Medtronic has more than 175,000 patients with implanted DBS systems worldwide and operates the largest dataset of stimulator-collected neural recordings ever assembled. Its public roadmap includes adaptive DBS for depression, the BrainSense ecosystem integration with consumer mobile devices, and a planned next-generation implant capable of bidirectional 64-channel recording and writing.

From a neuro-rights perspective, Medtronic represents the responsible-incumbent profile: a long compliance history under medical-device regulation, strong consent practices in clinical

use, but very limited disclosure about what data flows from implant to manufacturer servers and for what secondary uses. The black-box logging proposal advanced in later chapters of this book would have its largest near-term effect on Medtronic and a small number of comparable incumbents, because they have the most devices in the most heads.

Boston Scientific

Boston Scientific, founded in 1979, supplies the Vercise DBS system, which competes with Medtronic's Percept in the Parkinson's market. The Vercise Genus introduced multiple-independent-current-control programming for finer spatial selectivity. The company also markets the Precision SCS spinal-cord stimulator family and is investing in closed-loop pain management with sensing capability. Boston Scientific's recent acquisitions in the closed-loop epilepsy space and its partnership with NeuroPace bring its capabilities closer to those of Medtronic. The company is publicly committed to the proposed FDA Real World Evidence pathway, which could substantially shorten time-to-market for next-generation devices but also concentrates more neural data in manufacturer hands.

NeuroPace

NeuroPace, founded in 1997 in California, is a specialised maker of responsive neurostimulation systems for epilepsy. Its RNS System detects abnormal electrical activity in the brain and delivers brief stimulation pulses to interrupt seizures before the patient experiences them. As of 2026, the RNS has been implanted in more than 5,000 patients and has produced one of the largest databases of chronic intracranial EEG ever assembled. NeuroPace publishes physician-facing data dashboards but does not release patient-level neural recordings to third parties. Its planned expansion into responsive stimulation for depression and post-traumatic stress disorder is at clinical-trial stage.

Blackrock Neurotech

Blackrock Neurotech (formerly Blackrock Microsystems), based in Salt Lake City, is the long-standing supplier of research-grade intracortical microelectrode arrays. The Utah Array — a 96-channel silicon-based microelectrode array that has been the workhorse of academic BCI research for two decades — is a Blackrock product. The company markets the MoveAgain and NeuroPort systems for clinical-trial use in tetraplegic patients. Blackrock's MoveAgain has enabled paralysed individuals to control robotic arms, type on virtual keyboards, and operate digital cursors with thought alone. The company's planned next-generation array, Neuralace, is a higher-density wireless platform aimed at commercial deployment. Blackrock is unusual in that its devices are explicitly research-grade and require institutional review board oversight for every implantation; the regulatory transition to broad commercial use is the company's principal strategic challenge.

Synchron

Synchron, founded in 2012 and headquartered in New York, has developed the Stentrode, an endovascular BCI that is delivered to the motor cortex through a blood vessel rather than via open-skull craniotomy. The advantage is dramatic reduction in surgical risk and a shorter recovery period. Patients with the Stentrode have demonstrated text composition, web browsing, and home-device control with intention alone. Synchron's COMMAND clinical trial is in late stage and the company has received Breakthrough Device designation from the FDA. The Stentrode's lower surgical barrier could make it the first cortical BCI to reach broader clinical deployment, with all the regulatory and rights implications that brings. Synchron has publicly partnered with Apple to support a planned brain-control accessibility input protocol for iOS, which could be the moment commercial BCI capability ships to mass-market consumer devices.

Neuralink

Neuralink, founded in 2016 and headquartered in Fremont, California, is the most publicly visible and most controversial implant developer. Its N1 implant is a 1024-channel, fully wireless cortical recording device implanted by a custom surgical robot. As of 2026, Neuralink reports a single-digit number of human implants with limited functional demonstrations: cursor control, simple games, and basic text input. The company's public roadmap is ambitious — vision restoration through cortical writing, motor restoration through spinal-cord BCI, and eventually high-bandwidth general-purpose BCI for cognitively healthy users — but actual technical demonstrations remain modest relative to claims. The company's principal contribution may end up being its drive to standardise implant procedures and to demonstrate that a thousand-channel cortical device can operate reliably in a human. Its principal liability from a neuro-rights perspective is the same as its principal asset: visibility. Every regulatory and ethical question about cortical implantation now runs through Neuralink first.

Precision Neuroscience

Precision Neuroscience, founded in 2021 by former Neuralink co-founder Benjamin Rapoport, has developed the Layer 7 Cortical Interface — a flexible thin-film electrode array placed on the surface of the cortex through a small slot rather than penetrating the brain tissue. The reversibility of placement (the device can be removed without the kind of tissue damage that penetrating microelectrodes can cause) and the high channel count are the principal selling points. Precision received Breakthrough Device designation in 2024 and has completed early human studies.

Paradromics

Paradromics, founded in 2015 in Austin, Texas, has developed the Connexus Direct Data Interface, a high-bandwidth cortical implant aimed at communication restoration for severe paralysis. The device has been demonstrated in animal models and is in early human deployment. Paradromics emphasises high channel count and rapid data transmission, with the explicit goal of enabling speech decoding at usable rates.

Inbrain Neuroelectronics

Inbrain Neuroelectronics, founded in 2019 in Barcelona, has developed a graphene-based neural interface that promises higher charge-injection capacity than conventional electrodes, which in principle enables more selective stimulation with less tissue damage. The company has demonstrated the device in early human studies for Parkinson's disease.

Onward Medical

Onward Medical, headquartered in Eindhoven, Netherlands, develops spinal-cord stimulation systems for restoration of movement after spinal-cord injury and for the management of blood-pressure regulation in tetraplegic patients. Its ARC-EX (external) and ARC-IM (implanted) systems have demonstrated remarkable functional restoration in clinical trials, including supported walking in patients with complete spinal-cord injury. Although strictly speaking a spinal rather than brain device, Onward sits in the neurotechnology category because it interacts with the central nervous system.

Tier 2 — Clinical-grade Non-invasive Neurotechnology

g.tec medical engineering

g.tec, founded in 1999 in Austria, is the leading supplier of research-grade EEG and BCI systems to clinical and academic users worldwide. Its g.USBamp and g.HIamp amplifiers, the g.Nautilus wireless EEG system, and the g.tec Recoverix and mindBeagle clinical BCI systems for stroke rehabilitation and disorder-of-consciousness assessment are deployed in hundreds of clinical settings. g.tec does not sell to consumers; its devices are explicitly medical or research-grade.

Kernel

Kernel, founded in 2016 by Bryan Johnson, has developed the Kernel Flow, a time-domain near-infrared spectroscopy (TD-fNIRS) helmet that measures brain hemodynamics non-invasively at higher temporal resolution than conventional fNIRS. Kernel has shifted strategy several times — initially aiming at consumer markets, then pivoting to pharmaceutical research support — but the Flow remains one of the most sophisticated non-invasive neuroimaging devices commercially available.

Openwater

Openwater, founded in 2016 by Mary Lou Jepsen, is developing wearable optical imaging devices that aim to deliver imaging resolution approaching MRI from a much smaller form factor. The technology is at early demonstration stage but represents a category of capability — high-resolution non-invasive imaging — that, if realised, would substantially change the privacy calculus for neurotechnology.

Tier 3 — Consumer Wearable Neurotechnology

The consumer tier is where the largest population of users sits, where the regulatory framework is weakest, and where the data flows are most opaque. The Neurorights Foundation's 2024 review of thirty consumer neurotechnology companies found that twenty-nine of the thirty had access to users' brain data, imposed essentially no meaningful limits on that access, and nearly all reserved the right to share data with third parties. The companies below are a representative sample.

NeuroSky

NeuroSky, founded in 1999, is the original mass-market consumer EEG company. Its ThinkGear ASIC chip has powered dozens of toys, headbands and gaming peripherals over two decades. The company's footprint has shrunk as competitors entered, but its chip continues to power many low-end devices.

InteraXon (Muse)

InteraXon, founded in 2007 in Toronto, makes the Muse headband, a four-channel consumer EEG device targeted at meditation and sleep tracking. The Muse has been the best-selling consumer EEG device of the past decade. Its data-handling practices are typical of the category: broad licence in the terms of service, opt-in for research participation, no public commitment to specific limits on inference generation or third-party sharing.

Emotiv

Emotiv, founded in 2011 and based in San Francisco, supplies the Insight, EPOC X and MN8 EEG headsets to consumer, research and (increasingly) workplace markets. Emotiv is the company at the centre of the 2023 Chilean Constitutional Court ruling described in Chapter 8, which ordered the company to delete the brain data of plaintiff Guido Girardi. The case set the world's first court precedent on consumer neurodata and has reshaped Emotiv's terms of service for users in Chile.

BrainCo

BrainCo, founded in 2015 by Harvard-affiliated researchers, makes the Focus 1 and FocusEDU headbands marketed to schools and to consumer focus-enhancement applications. The deployment of FocusEDU in Chinese classrooms has been one of the most cited examples of school neuro-surveillance worldwide and has provoked significant public debate.

Neurable

Neurable, founded in 2015 in Boston, makes the MW75 Neuro headphone, a consumer audio product with embedded EEG sensors that purport to track attention and focus. The integration of EEG into an existing consumer-electronics form factor — headphones — is a template that more consumer-electronics manufacturers are likely to follow.

Cognixion

Cognixion, founded in 2014, makes the Cognixion ONE, an augmented-reality headset with embedded EEG sensing aimed at people with severe motor disability who cannot speak. The device combines a head-mounted display, EEG-based intention detection and AI-driven communication. Its medical/non-medical positioning is intermediate: it is marketed to clinical users but not yet FDA-cleared as a medical device in all use cases.

Mendi, Neurosity, Bitbrain, Neeuro and the long tail

A long tail of smaller consumer companies — Mendi (Sweden, fNIRS headbands), Neurosity (US, dry-electrode EEG for software developers), Bitbrain (Spain, research-to-consumer EEG), Neeuro (Singapore, education-focused EEG), and many others — populates the lower-cost end of the market. The category is fragmenting rather than consolidating, which complicates regulatory enforcement: a regulator who can monitor a handful of major firms struggles to keep up with dozens of small ones.

Adjacent Players — The AI and Platform Layer

No survey of the neurotechnology industry would be complete without acknowledging the role of the AI and platform companies that increasingly sit between the device maker and the data. Apple, with its Brain-Computer Interface Human Interface Device input protocol announced in 2025, is staking out a position as the standardiser of BCI input across consumer operating systems. Google DeepMind has published research on foundation models for neural decoding. Meta operates Reality Labs, whose acquisition of CTRL-Labs in 2019 gave it a substantial position in EMG-based intention detection. Microsoft Research has neural-interface programmes that, while not productised, contribute to the academic state of the art. These platform companies have no implants in heads today, but they will set the protocols by which most implant data is processed tomorrow.

OpenAI, Anthropic and other foundation-model companies are also relevant, even though they do not produce neural interfaces, because their models will increasingly be the inference layer that turns raw brain signals into actionable inferences. A regulatory regime that targets the device makers but leaves the inference companies untouched will miss most of the risk.

What This Industry Picture Means for Neuro-rights

Three policy conclusions flow from the atlas. First, the regulatory regime must treat the device makers as a heterogeneous category. A medical-device regulator that knows how to license a DBS implant has no template for licensing a meditation headband; a consumer-protection regulator that knows how to police misleading marketing has no template for auditing AI inferences on neural data. Each tier requires its own competent authority, and the proposal in Chapter 67-bis to create a National Neurotechnology Commission with cross-tier jurisdiction follows from this observation.

Second, the value capture is migrating. The early hardware era is yielding to a data and software era; the company that sells the implant is rarely the one that will profit most from the data it generates. Apple, Google, Meta and the foundation-model companies are positioning themselves to be the principal beneficiaries of the platform layer. A regulatory regime that targets implant makers but leaves platforms untouched will protect the wrong actors at the wrong layer.

Third, the smaller companies matter more than their balance sheets suggest. The neuro-rights catastrophes most likely to actually happen — schoolchildren tracked, employees monitored, vulnerable patients exposed — will most likely involve consumer wearables and low-cost workplace devices rather than billion-dollar implants. A regulator that focuses

entirely on Neuralink and Synchron will miss the actual locus of harm. The chapter on agencies (Chapter 70-bis) develops the implications.

CHAPTER 63

The Military-Industrial-Neuro Complex

If the commercial neurotechnology industry represents the visible tip of a rapidly growing iceberg, the military sector constitutes the vast, submerged mass beneath the surface—far larger, far better funded, and far less transparent. Understanding the full landscape of neurotechnology development requires acknowledging a fundamental reality: the most advanced research into the human brain is not being conducted by Silicon Valley startups or medical device companies, but by defense agencies and their contractors. This is the military-industrial-neuro complex, and it has been driving neurotechnology R&D for decades.

DARPA: The Engine of Military Neuroscience

The Defense Advanced Research Projects Agency (DARPA) is the single most important institution in the history of neurotechnology. As the agency itself stated in its 2003 Strategic Plan, the "long-term Defense implications of finding ways to turn thoughts into acts, if it can be developed, are enormous: Imagine U.S. warfighters that only need use the power of their thoughts to do things at great distances." This vision—controlling weapons with the mind—has driven DARPA's neuroscience investments for over two decades.

DARPA's track record lends credibility to even its most ambitious goals. Among the agency's past accomplishments are the Saturn rocket, ground radar, the Stealth Fighter, the Predator missile, and the Internet itself. With a fiscal year 2005 budget of $3 billion (a figure that has only grown since), DARPA has the resources and patience to pursue long-term, high-risk projects that no commercial entity would attempt.

The agency's neuroscience portfolio has been remarkably broad. Its Continuous Assisted Performance (CAP) program sought to "investigate ways to prevent fatigue and enable soldiers to stay awake, alert, and effective for up to seven days straight without suffering any deleterious mental or physical effects and without using any of the current generation of stimulants." A project called "Metabolic Dominance" aimed to develop a "nutraceutical" that would vastly improve soldiers' endurance. The agency has also invested in sleep-reduction experiments and research into whether the body's core temperature can be altered depending on weather conditions. The ultimate goal, DARPA stated, is "to enable superior physical and physiological performance by controlling energy metabolism on demand."

From Monkeys to Battlefields

DARPA's brain-machine interfaces program achieved a milestone in the early 2000s that would have seemed like science fiction a generation earlier. Researchers demonstrated that a monkey could control a robot arm using only neural impulses. In a 2002 report to Congress, DARPA insisted that these were not merely the impulses that normally control the monkey's

arm, but “the very thoughts about arm movement themselves, transmitted to a robot limb in another room.”

At Duke University, scientists trained monkeys to engage in complicated movements using neural processes alone, involving both reaching for and grasping an object. At the 2003 meeting of the Society for Neuroscience, Duke researchers noted they had proved the same principle in humans undergoing neurosurgery. The implications were clear: once neural impulses could be captured and digitized, they could be sent as encrypted messages over the Internet to “do things at any distance the electrons will travel, including maneuvering aircraft, inspecting a target, releasing weapons, and so forth.”

The Broader Military Ecosystem

DARPA is not alone. The Intelligence Advanced Research Projects Activity (IARPA) has pursued parallel efforts, as documented in the 2021 Naval Postgraduate School thesis “Neurowar Is Here!” by Lieutenant Colonel Joshua D. Gramm and Major Brian A. Branagan. The U.S. Air Force, Army, Navy, Marine Corps, and Special Operations Command all maintain active research programs in neurotechnology.

The Army Research Laboratory, for example, has explored the use of directed energy systems including the Active Denial System—a millimeter-wave weapon that causes an intense heating sensation on the skin. The Air Force Office of Scientific Research has funded work on cervical transcutaneous vagal nerve stimulation (ctVNS) to enhance cognitive performance in pilots. The Defense Health Agency’s Center for Military Psychiatry and Neuroscience studies the effects of combat on brain function.

This ecosystem operates through a complex web of contracts, Cooperative Research and Development Agreements (CRADAs), and classified programs. The privatization of brain research means that defense contractors—Lockheed Martin, Raytheon, Northrop Grumman, and others—are deeply embedded in the development of military neurotechnology.

The Ethical Vacuum in Military Research

The ethical questions raised by this military neuroscience enterprise are profound. As bioethicist Jonathan Moreno wrote in his 2004 article “DARPA on Your Mind”: “How will individuals be recruited to try what will obviously be very dangerous experiments? If such experiments are done in secret, how will they be ethically reviewed? What are the dangers of such information falling into ‘the wrong hands,’ and are there any ‘right hands’ for this kind of knowledge?”

Moreno observed that “discussions about ethics require the oxygen of transparency, precisely the item in short supply in national security matters.” He proposed two remedies: a funded research program to examine the ethical issues publicly, and the creation of an “ethical, legal and social implications” advisory panel within agencies like DARPA, composed of individuals with appropriate security clearances but who are not federal employees.

Two decades later, neither remedy has been fully implemented. The military-industrial-neuro complex continues to operate with far less public scrutiny than the commercial

neurotechnology sector—and with far greater resources. As the UN report on neurotechnology and human rights noted, the very technologies that could enhance soldiers could also be weaponized against civilians. The line between defense and offense, between protection and control, has never been clear.

The Shadow of History

The current military neuroscience enterprise does not exist in a historical vacuum. As Moreno documented, the U.S. government has a long history of secret brain research. During the 1950s, the CIA's MKULTRA program dosed unsuspecting individuals with LSD, including Army anthrax expert Frank Olson, who fell to his death from a New York City hotel in 1953. The Army Chemical Corps conducted involuntary experiments with mescaline at New York State's Psychiatric Institute. Thousands of soldiers were given LSD in tests to which their consent was questionable.

These experiments occurred despite the Nuremberg Code's first principle: "The voluntary consent of the human subject is absolutely essential." Less than two months after one victim's death, the Secretary of Defense issued a top-secret memorandum making the Nuremberg Code the Pentagon's policy for atomic, biological, and chemical warfare experiments. Yet the violations continued.

The lesson is not that today's military researchers are repeating past abuses. It is that secrecy, national security imperatives, and the absence of independent oversight create conditions where abuse becomes possible. The military-industrial-neuro complex, with its classified programs and its pursuit of ever more powerful tools to enhance and control the human brain, operates in exactly such conditions today.

CHAPTER 64

The Investment Ecosystem — Who Funds the Future of the Brain?

Behind every neurotechnology breakthrough is a chain of investment decisions. Understanding the structure of neurotechnology investment — who funds what, with what expectations, on what timescales — is essential to understanding the trajectory of the industry and the opportunities for shaping its development.

The Scale of Investment

Private investment in neurotechnology grew from $331 million in 2014 to $7.3 billion in 2023 — a more than twenty-fold increase in less than a decade. The figure understates true investment, as it excludes government funding, military procurement, and the substantial in-house research budgets of major technology companies that do not appear in venture capital data.

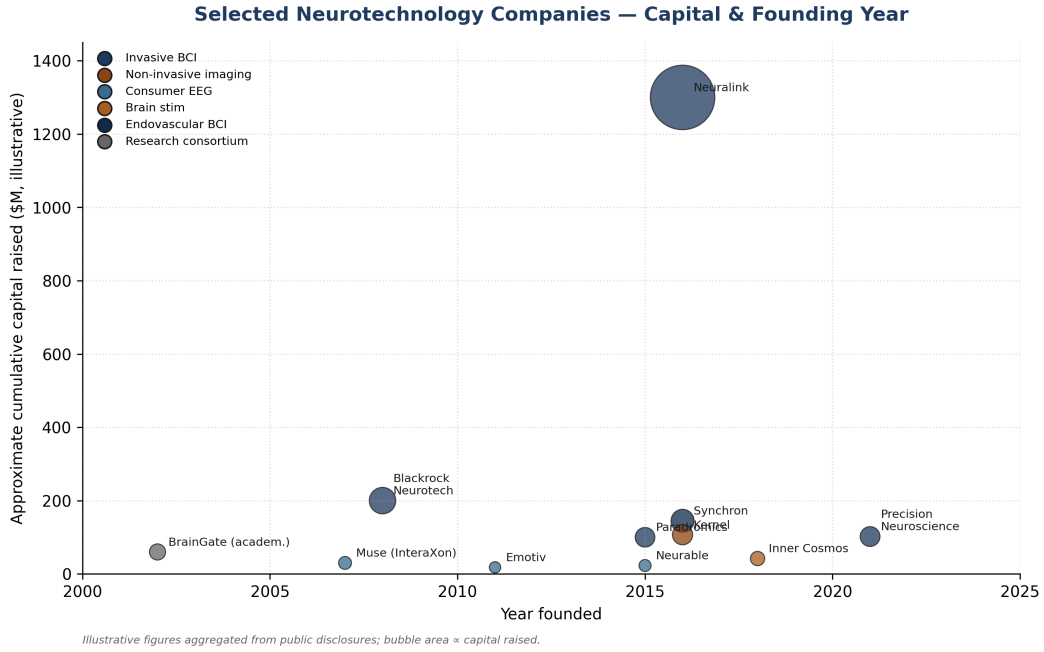


Figure 45.1. *Selected neurotechnology companies by founding year and approximate cumulative capital raised. Bubble area is proportional to capital. Figures are illustrative and aggregated from public disclosures.*

When public investment is added, total funding for neurotechnology research and development exceeds $25 billion annually worldwide. Major national programmes include the United States BRAIN Initiative ($6 billion cumulative through 2024), the European Human Brain Project (€1 billion), China's China Brain Project ($10 billion projected over 15 years), and Japan's Brain/MINDS programme. These programmes fund both basic neuroscience and applied technology development.

The Venture Capital Landscape

A small number of venture capital firms dominate neurotechnology investment. Founders Fund (Peter Thiel), Khosla Ventures, GV (Google Ventures), Bezos Expeditions, and a handful of others account for a substantial fraction of total private investment. These firms typically focus on companies with potential for massive scale — brain-computer interfaces with consumer applications, decoding platforms with broad licensing potential, and pharmaceutical-class therapeutics.

The investment logic mirrors that of other deep-tech sectors: long development timelines (typically ten to twenty years), high capital requirements (often $500 million or more before profitability), high regulatory risk, and — if successful — returns that justify the risk through near-monopoly positions in markets of enormous scale. Neuralink, valued at $5 billion in 2023, exemplifies the pattern: years of pre-revenue development, sustained funding from a small group of high-conviction investors, and a market opportunity that, if realised, justifies the capital deployed.

The Big Tech Position

Major technology companies have taken substantial positions in neurotechnology, both directly and indirectly. Meta's acquisition of CTRL-Labs in 2019 for an estimated $1 billion brought neural interface capability in-house at one of the world's largest consumer technology companies. Google's parent Alphabet has invested in multiple neurotechnology companies through GV and has internal research programmes in neural interfaces. Apple has filed substantial patents in biosignal processing and is reported to have internal neural interface programmes. Microsoft has invested in brain-computer interface research through several initiatives.

The strategic logic is straightforward: brain-computer interfaces represent a potential next platform shift, comparable to the mobile platform shift of the 2000s. Whoever controls the dominant neural interface platform will, on plausible projections, capture economic value comparable to the dominant smartphone platforms of the present. The investment is justified by the magnitude of the platform opportunity even if probability of success is modest.

Concentration Risk

The neurotechnology investment landscape exhibits exceptional concentration. A small number of investors, companies, and technical approaches dominate. This concentration creates both opportunity and risk: opportunity, because targeted policy can shape the industry by engaging a manageable number of actors; risk, because the technology trajectory is determined by the priorities of a narrow elite whose interests may not align with broader social interests.

Government and Defence Investment

Beyond the public neurotechnology programmes mentioned above, substantial defence investment flows into neurotechnology in major military powers. In the United States, DARPA programmes including N3 (Next-Generation Non-Surgical Neurotechnology), RAM (Restoring Active Memory), and NESD (Neural Engineering System Design) have funded foundational neural interface research with explicit dual-use objectives. The Office of Naval Research, the Air Force Research Laboratory, and the Army Research Laboratory all maintain neurotechnology programmes.

Comparable programmes exist in China, Russia, and other military powers, though the level of transparency varies substantially. The Chinese People's Liberation Army has been explicit about its interest in cognitive warfare and military brain science. Russian research traditions in psychotronics extend back to the Soviet era and have been documented in academic literature.

The intersection of civilian and military neurotechnology investment is significant. Companies funded primarily by venture capital may receive secondary funding from defence agencies; technologies developed under defence funding may be commercialised through civilian companies. The boundary between civilian and military neurotechnology is permeable, with implications for export controls, dual-use regulation, and the application of laws of armed conflict.

The Pension Fund Question

A quietly important development is the increasing position of public pension funds in neurotechnology investment. CalPERS, the California Public Employees' Retirement System, and similar large public pension funds invest substantially in venture capital funds that hold neurotechnology positions. State teachers' retirement systems, sovereign wealth funds, and public university endowments are similarly exposed.

The implication is that public-sector entities — entities theoretically accountable to democratic processes — are funding the development of technologies whose deployment they have no specific authority to approve or constrain. The investment decisions are made by professional fund managers operating under fiduciary duties to maximise returns. Whether and how the public interest in neurotechnology development is reflected in these investment decisions is a question that pension fund governance has not adequately addressed.

Ethical Investment and the Neuro-rights Movement

A nascent ethical investment movement is beginning to address neurotechnology. The Neurorights Foundation has engaged with institutional investors on neural-data privacy practices. Several activist investors have raised neuro-rights concerns at shareholder meetings of major technology companies. ESG (Environmental, Social, Governance) rating agencies are beginning to incorporate neurotechnology concerns into their assessments.

The movement is in its early stages. Compared to the ethical investment movement around fossil fuels, weapons, or tobacco, neurotechnology ethical investment has not yet achieved

comparable scale or institutional traction. But the foundations are being laid for a broader engagement that could, over time, shape the cost of capital for companies developing neurotechnology in ways that ignore human rights considerations.

Implications for Policy

The structure of neurotechnology investment has several implications for policy:

38. Engagement with a small number of high-impact investors can shape the industry more efficiently than diffuse regulation. A handful of venture capital firms determine which technologies receive the funding required to scale.
39. Big Tech's position requires specific attention. Antitrust enforcement, platform regulation, and acquisition review must address neurotechnology positions of major technology firms.
40. The dual-use boundary requires specific export control and procurement scrutiny. Technologies developed for civilian markets with defence applications, and vice versa, must be subject to coherent oversight.
41. Pension fund governance must address neurotechnology positions. Public pension funds invested in companies developing dual-use neurotechnology have governance obligations that current frameworks do not adequately specify.
42. Ethical investment infrastructure must mature. ESG frameworks, shareholder engagement standards, and divestment criteria for neurotechnology must be developed to support the emerging ethical investment movement.

PART XII

The Legislative Agenda

The existing legal frameworks for neuro-rights, while groundbreaking in their ambition, leave critical gaps that render them insufficient to protect citizens from the full spectrum of neurotechnological harm. The UNESCO Recommendation on the Ethics of Neurotechnology, adopted in November 2025 by more than 190 member states, provides the first global ethical framework—but it is soft law, a recommendation without binding force. Chile's constitutional reform, the EU AI Act, Colorado and California's neurodata protections—each advances the cause, yet each leaves entire categories of abuse unaddressed.

The legislative agenda that follows emerges from a careful mapping of what exists against what is needed. It draws on the documented experiences of 248 signatories from 36 countries who have reported consistent patterns of non-consensual technological targeting, and on the technical and legal analysis developed across the preceding sections of this guide.

The Legislative Gap, Defined

Three structural deficiencies recur across every jurisdiction examined:

First, existing statutes target consumer device-makers. The GDPR, the AI Act, and state-level neurodata laws were written with commercial brain-reading headsets and wellness apps in mind. They are silent on the actors with the strongest incentive to read minds without genuine consent—the state, the employer, the school, and the insurer. A company that markets an EEG headband must comply with data protection rules; an intelligence agency that deploys remote neural monitoring faces no equivalent constraint.

Second, no binding instrument addresses weaponisation. The Chemical Weapons Convention and Biological Weapons Convention prohibit entire classes of weapons. No

comparable treaty exists for neuroweapons—directed-energy devices, microwave and ultrasonic systems, and other technologies capable of acting on the human nervous system at a distance. The result is a legal vacuum in which the very capabilities documented in DARPA's N3 program and described in the Hi-Tech War Crime Report can be developed, tested, and potentially deployed against civilians without violating any international law specific to their use.

Third, no mechanism exists for the individual victim. A person who believes they are being subjected to non-consensual neurotechnological targeting has no dedicated offence to cite, no defined complaint procedure, and no technically competent body to investigate. As the TI One Voice petition to the European Parliament documents, those who come forward are "too often turned away at the door—not because their claims have been examined and found wanting, but because no mechanism exists to examine them at all."

The Single Highest-Value Legislative Fix

Among the many possible reforms, one stands out as disproportionately consequential: the definition of "neural data." US states have already split on exactly this point. Colorado's privacy law protects raw neural signals; California's goes further, protecting the algorithmic inferences drawn from those signals. The difference matters enormously. If the law protects only the raw signal, a company can collect brain data, run it through an AI decoder, and claim the resulting mental states are not "neural data" but proprietary analysis. The inference, not the signal, is what reveals a person's thoughts.

Any neuro-rights legislation at any level—national, regional, or international—must define neural data to include both the raw signal and the inferences derived from it. This single clause determines whether the law has teeth or is a paper shield.

The Missing Offences

No criminal code in any jurisdiction currently includes an offence of "non-consensual neurotechnological interference." The proposed legislative reform packages developed by TI One Voice for the Belgian Federal Parliament, the European Parliament, and the US Congress address this directly. The Belgian proposal, for example, introduces a new Article 422quater covering the deliberate use of electromagnetic, acoustic, or other directed-energy devices to cause physical harm, psychological distress, sleep deprivation, or cognitive interference. The offence applies "whether or not such use leaves externally visible injuries"—a critical provision given that neuroweapons may produce no visible wounds while causing profound neurological damage.

The same reform packages propose expanding stalking provisions to cover coordinated multi-actor harassment and technology-mediated harassment including GPS tracking, drone surveillance, and acoustic harassment devices. They also call for a statutory duty on police to log and aggregate complaints of coordinated harassment patterns, even where individual incidents appear minor in isolation.

The Institutional Gap

Even with perfect laws, enforcement requires institutional capacity. No existing oversight body has the technical expertise to audit electromagnetic emissions from suspected government installations, to distinguish between a psychiatric condition and the effects of directed energy, or to investigate a complaint of remote neural monitoring.

The proposed reforms address this through three mechanisms: specialised technical assessment units within intelligence oversight bodies, independent diagnostic centres equipped with RF spectrum analysis and neuroimaging capabilities, and parliamentary inquiry commissions with binding investigative powers. The Belgian proposal models such a commission on the precedent of the Dutroux Commission of 1996—a demonstration that when a pattern of abuse is sufficiently documented, democratic institutions can and do respond with serious investigative machinery.

The Victim Recognition Gap

Perhaps the most fundamental gap is the absence of recognised victim status. The Victims' Rights Directive of the European Union provides a framework for support, but it does not contemplate technological targeting. A person subjected to sustained neuroweapons harassment may meet the directive's definition of a victim—they have suffered harm—butbut the systems designed to help them have no category for the harm they have experienced. Medical professionals misdiagnose technological torture as mental illness. Police lack protocols for documenting invisible injuries. The legal system offers no remedy because the wrong has no name.

The TI One Voice petition asks the European Parliament to recognise technological targeting within the Union's victims'-rights framework, ensuring access to competent, impartial medical and technical evaluation. This is not a request for special treatment. It is a request for equal treatment—for the same right to have one's complaint investigated that any citizen of a democratic state should expect.

A Mandatory Black Box

One achievable measure would transform the landscape of accountability: a mandatory tamper-evident logging requirement for neural devices. Every commercial or government neurotechnology capable of recording or stimulating neural activity should contain a "black box" that records its own operation—when it transmitted, what it transmitted, and to whom. Such a requirement would make abuse provable after the fact. It would deter misuse by removing the plausible deniability that currently protects perpetrators.

No existing legislation requires such logging. The AI Act does not mention it. The GDPR does not require it. The UNESCO Recommendation recommends transparency but does not mandate technical mechanisms. This is a gap that can be closed with a single legislative clause, and it would do more to protect citizens than any number of ethical principles.

The Legislative Agenda in Summary

The missing legislation falls into five categories:

43. **A definition of neural data** that includes raw signals and algorithmic inferences, applied uniformly across all contexts—consumer, medical, employment, and state use.
44. **Criminal offences** for non-consensual neurotechnological interference, directed-energy assault, and technology-mediated harassment, with penalties that reflect the gravity of the harm and aggravated circumstances for state actors.
45. **An international treaty** banning neuroweapons, modelled on the Chemical Weapons Convention and Biological Weapons Convention, with verification and enforcement mechanisms.
46. **Institutional capacity** for investigation: technical audit units within oversight bodies, independent diagnostic centres, and parliamentary inquiry mechanisms with binding powers.
47. **Victim recognition and support**: inclusion of technological targeting within victims'-rights frameworks, healthcare protocols that distinguish between technological harm and psychiatric conditions, and specialised support services.

Each of these can be pursued at multiple levels simultaneously—constitutional amendments at the national level, regulations at the regional level, and treaties at the international level. The UNESCO Recommendation provides the ethical foundation. The missing legislation provides the legal architecture. What remains is the political will to build it.

CHAPTER 65

The Missing Legal Framework — Five Critical Gaps

The UNESCO Recommendation of 2025 provides the ethical foundation. It declares that neurotechnology must respect human dignity, that mental privacy must be protected, and that cognitive liberty is a fundamental value. But a recommendation is not a law. It has no enforcement mechanism, no penalties, no investigative capacity, and no victim compensation fund.

Between the ethical principle and the enforceable right lies legislation. And legislation, in the domain of neurotechnology, is almost entirely absent.

This chapter identifies the five most critical gaps in current law. Each is a hole through which human rights violations can pass without detection, without remedy, and without accountability.

Gap One: No Definition of Neuro-Crimes

No criminal code in any country contains an offence called "neurotechnological assault," "cognitive interference," or "non-consensual neural manipulation." The closest existing provisions are assault statutes, which require physical contact or the threat of it, and stalking statutes, which require patterns of behaviour that are visible to conventional investigation.

Directed-energy weapons that cause neurological harm without leaving externally visible injuries fall into a legal void. The victim experiences tinnitus, sleep disruption, cognitive fog, and unexplained physiological effects—symptoms that the criminal law does not recognise as evidence of a crime. As the OHCHR submission from TI One Voice documents, "victims lack formal complaint procedures that recognise technological attacks" and "medical professionals misdiagnose technological torture as mental illness."

The proposed amendment to the Belgian Criminal Code—Article 422quater—offers a model. It would create a specific offence covering "the deliberate use of electromagnetic, acoustic, or other directed-energy devices to cause physical harm, psychological distress, sleep deprivation, or cognitive interference to a person, whether or not such use leaves externally visible injuries." It includes aggravating circumstances for state actors and for sustained campaigns. It reverses the evidentiary burden regarding technical feasibility once the victim demonstrates a pattern of symptoms consistent with published bioeffects literature.

No jurisdiction has yet enacted such a provision. This is the first gap.

Gap Two: No Comprehensive Neurodata Protection

The GDPR protects personal data. The AI Act protects against certain forms of algorithmic manipulation. The Colorado and California privacy laws protect neural data as sensitive information. But none of these instruments defines "neural data" to include the *algorithmic inferences drawn from it*, and none applies uniformly across all contexts.

The consequence is a patchwork that leaves the most vulnerable unprotected. As the EU neuro-rights guide notes, “existing statutes target consumer device-makers and are largely silent on the actors with the strongest incentive to read minds without genuine consent—the state, the employer, the school and the insurer.”

A comprehensive neurodata protection framework would:

Define neural data to include raw signals and all inferences derived from them

Apply to all actors: commercial, governmental, medical, and educational

Require explicit, informed, and revocable consent for any collection or processing

Prohibit the use of neural data for discrimination in employment, insurance, education, and criminal justice

Mandate tamper-evident logging for all neural devices, creating an auditable record of when data was collected, by whom, and for what purpose

No such framework exists anywhere in the world. This is the second gap.

Gap Three: No Neuroweapons Ban

The Chemical Weapons Convention bans the development, production, stockpiling, and use of chemical weapons. The Biological Weapons Convention does the same for biological weapons. There is no equivalent for neuroweapons.

The technologies described in this guide—directed-energy devices capable of inducing neurological effects, brain-computer interfaces that can be weaponised, AI systems that can coordinate cognitive warfare—are not prohibited by any binding international instrument. The UNESCO Recommendation is soft law. The AI Act addresses software but not hardware. The Council of Europe Framework Convention on AI covers artificial intelligence but not the directed-energy systems that can be used independently of AI.

The result is a regulatory vacuum in which neuroweapons can be developed, tested, and deployed without violating any treaty obligation. As the legislative reform proposals from TI One Voice argue, what is needed is “an absolute prohibition on the use of directed-energy weapons and neuro-technological devices against any person within EU territory, except under strictly defined law-enforcement scenarios subject to judicial authorisation.”

At the international level, the model exists: the Chemical Weapons Convention provides a template for verification, enforcement, and victim assistance that could be adapted to neuroweapons. But no state has yet proposed such a treaty. This is the third gap.

Gap Four: No Oversight Mechanism

Even where laws exist, they are only as effective as the institutions that enforce them. No country has established an independent body with the technical expertise to audit neurotechnology, investigate complaints of misuse, and hold perpetrators accountable.

The Belgian intelligence oversight body (Comité R) has no mandate to audit directed-energy capabilities. The European Commission has no neurotechnology unit. The UN Human

Rights Council has no special rapporteur on neurotechnological violations. The OHCHR submission calls for the establishment of exactly such a mechanism, noting that "investigative bodies lack technical expertise to verify neurotechnological harassment."

The legislative reform proposals outline what oversight should look like:

A specialised technical assessment unit within existing oversight bodies, staffed with engineers and physicists with appropriate security clearances

Mandatory annual public reporting on the acquisition, possession, and deployment of neurotechnology capabilities by state actors

A citizen complaint pathway with statutory response deadlines

Independent diagnostic centres equipped with RF spectrum analysis, audiological assessment, and neuroimaging capabilities

None of these institutions exists. This is the fourth gap.

Gap Five: No Victim Recognition

Perhaps the most fundamental gap is the absence of legal recognition for victims of neurotechnology abuse. Without recognition, there is no pathway to justice, no access to support services, and no official record that the harm occurred.

The OHCHR submission describes the consequences with brutal clarity: "Victims lack formal complaint procedures that recognise technological attacks. Investigative bodies lack technical expertise to verify neurotechnological harassment. Medical professionals misdiagnose technological torture as mental illness."

The Targeted Individual phenomenon—thousands of people worldwide reporting sustained technological harassment—remains officially unrecognised. Their complaints are dismissed as delusional. Their symptoms are treated as psychiatric conditions. Their testimony is excluded from evidence because the legal system has no framework for evaluating it.

The legislative reform proposals address this gap by calling for:

Inclusion of TI status as a recognised category within victim-support frameworks

Healthcare protocols that require consideration of environmental and technological causes before issuing psychiatric diagnoses

Specialised diagnostic centres for independent medical-technical evaluation

Training for medical professionals to distinguish between psychological conditions and technological interference

No country has yet taken these steps. This is the fifth gap.

The Cost of Inaction

These five gaps are not abstract. They are the reason that victims of neurotechnology abuse currently have no remedy. They are the reason that perpetrators operate with impunity. They

are the reason that the technologies described in this guide can be developed, deployed, and weaponised without legal constraint.

Closing each gap requires legislation. The UNESCO Recommendation provides the ethical foundation. The missing legal framework provides the architecture. What remains is the political will to build it.

CHAPTER 66

At What Level Should Neuro-rights Be Legislated?

The question of where neuro-rights should be codified is not merely procedural. It is strategic. The level at which a right is protected determines its enforceability, its reach, and its vulnerability to political change. Legislators face a choice among four distinct tiers: constitutional amendment, national statute, regional regulation (such as an EU directive or regulation), and international treaty. Each tier carries specific advantages and limitations, and each interacts differently with the others.

The Constitutional Tier: Maximum Entrenchment, Maximum Difficulty

Chile demonstrated the constitutional path in 2021, amending its constitution to protect brain rights. The advantage of this approach is clear: constitutional rights are the most durable form of legal protection. They cannot be repealed by ordinary legislation, and they create a floor beneath which no statute may fall. Constitutional neuro-rights would bind all branches of government—executive, legislative, and judicial—and would provide a basis for constitutional review of any law or action that infringes on mental privacy, personal identity, or cognitive liberty.

The disadvantage is equally clear. Constitutional amendment requires supermajorities, often multiple legislative votes across successive sessions, and sometimes popular referenda. In politically divided systems, this threshold may be unreachable. Chile succeeded because of a unique moment of constitutional renewal following mass protests. Most countries lack such a window.

Moreover, constitutional rights are typically expressed at a high level of abstraction. “The right to mental privacy” must be given concrete content through implementing legislation. The constitution declares the principle; the statute defines the offence, establishes the enforcement mechanism, and sets the penalties. Constitutional neuro-rights without statutory implementation risk remaining aspirational.

The National Statutory Tier: Speed and Specificity

National legislation offers the fastest path to enforceable neuro-rights. A well-drafted statute can define neural data with precision, establish criminal offences for non-consensual neurotechnology use, create oversight mechanisms, and provide remedies for victims. The Colorado and California neurodata privacy laws of 2024 demonstrate what is possible at the state level within federal systems.

The legislative reform proposals developed by TI One Voice for the Belgian Federal Parliament illustrate the statutory approach in detail. They propose amendments to the Criminal Code creating a specific offence of “directed-energy and neuro-technological assault” (new Article 422quater), expanding stalking provisions to cover organised covert

harassment, and reforming intelligence oversight to require public reporting on neurotechnology capabilities. These measures can be enacted through ordinary legislative processes, without constitutional amendment.

The limitation of national legislation is jurisdictional. Neurotechnology does not respect borders. A device operated in one country can affect a person in another. National laws cannot regulate the conduct of foreign states or foreign corporations, and they cannot compel cooperation from foreign law enforcement. A patchwork of national statutes, however well-drafted, will leave gaps that abusers can exploit.

The Regional Tier: Scale and Harmonisation

Regional bodies like the European Union offer an intermediate tier that combines scale with relative legislative speed. An EU regulation is directly applicable in all 27 member states, creating uniform protection across a population of 450 million. An EU directive sets minimum standards that member states must transpose into national law, allowing for local adaptation while ensuring a common baseline.

The EU already has relevant instruments: the GDPR, the AI Act (Regulation 2024/1689), and the Charter of Fundamental Rights. But as the TI One Voice submission to the European Parliament notes, "the EU has no dedicated neuro-rights instrument; it relies on a patchwork of the GDPR, the AI Act and a 2019 OECD recommendation, none of which was written with brain-reading consumer devices in mind."

The legislative reform proposals for the European Parliament address this gap through three coordinated measures: a new Regulation on the Prohibition of Covert Directed-Energy and Neuro-Technological Interference Against Persons, amendments to the AI Act extending its prohibited practices to cover hardware-based neuro-interference, and a Directive on Victim Protection for Technology-Facilitated Covert Harassment. These would create a comprehensive regional framework that no single member state could achieve alone.

The limitation of the regional tier is that it cannot bind non-member states. A Chinese or Russian company developing neuroweapons would not be subject to EU law. Nor would the intelligence services of non-member states. Regional protection is powerful within its territory but has no extraterritorial reach against state actors.

The International Treaty Tier: Universality and Legitimacy

An international treaty banning non-consensual neurotechnology use—modelled on the Chemical Weapons Convention and the Biological Weapons Convention—would provide the highest level of legitimacy and the broadest jurisdictional reach. It would bind all states parties, create obligations to enact implementing legislation, establish verification mechanisms, and provide a framework for international cooperation in enforcement.

The UNESCO Recommendation on the Ethics of Neurotechnology, adopted in November 2025, provides the ethical foundation for such a treaty. But a recommendation is soft law. It has no binding force. The step from recommendation to treaty is the step from aspiration to obligation.

The arguments for a treaty are strong. Neuroweapons, like chemical and biological weapons, pose a threat to humanity as a whole. No single state can defend against them alone. A treaty would create a global norm against the weaponisation of the brain, making the use of such weapons a matter of international concern rather than a domestic policy choice.

The obstacles are equally strong. Treaty negotiation takes years, sometimes decades. Verification is extraordinarily difficult for technologies that may be indistinguishable from ordinary electromagnetic emissions. And the states most likely to develop neuroweapons are the least likely to sign a treaty restricting them.

The Principle of Subsidiarity

The European Union's principle of subsidiarity provides a useful framework for thinking about the optimal legislative level. Subsidiarity holds that decisions should be made at the lowest level of governance capable of achieving the objective, unless action at a higher level is clearly more effective.

Applied to neuro-rights, subsidiarity suggests a multi-level approach. Constitutional amendment should be pursued where possible, to provide the strongest possible foundation. National legislation should fill the gaps in criminal law, victim support, and healthcare protocols. Regional regulation should ensure harmonisation across borders and address the challenges of scale. International treaty should be pursued as a long-term objective, to establish global norms and provide a framework for cooperation.

The legislative reform proposals for Belgium, the European Parliament, and the United States Congress reflect this multi-level thinking. Each package is tailored to its jurisdiction's constitutional structure and existing legal instruments, while maintaining coherence across borders. The goal is not to choose one level over the others, but to create a system in which each level reinforces the others, and no gap remains through which perpetrators can escape accountability.

The Harmonisation Challenge

Multi-level legislation creates its own challenges. Inconsistencies between levels can create loopholes. A definition of "neural data" that differs between national and regional law can lead to forum shopping. A criminal offence that exists in one jurisdiction but not another can prevent extradition.

Harmonisation requires coordination. The TI One Voice proposals address this by using consistent definitions across all three legislative packages, by referencing common technical standards (IEEE C95.1, ICNIRP guidelines), and by proposing mechanisms for cross-border cooperation. The EU Directive on Victim Protection would require member states to designate competent authorities capable of cooperating with their counterparts in other jurisdictions.

The most consequential harmonisation question concerns the definition of neural data. As the TI One Voice submission to the European Parliament notes, "the single highest-value legislative fix is defining 'neural data' to include the algorithmic inferences drawn from it,

not only the raw signal. US states have already split on exactly this point, which makes the definition the most consequential clause in any bill."

If Colorado defines neural data one way and California defines it another, a company can structure its operations to avoid the stricter definition. If the EU defines it one way and the United States another, a global company can choose which market's standards to meet. Harmonisation at the definitional level is therefore a prerequisite for effective protection at any level.

The Practical Path Forward

The most realistic path is not to choose a single level but to advance on all levels simultaneously, recognising that different levels serve different functions. Constitutional amendment provides durability. National legislation provides specificity and speed. Regional regulation provides scale and harmonisation. International treaty provides universality and long-term legitimacy.

The urgent priority is national legislation to close the five critical gaps: the absence of definitions for neuro-crimes, the lack of neural data protection, the absence of a neuroweapons ban, the lack of oversight mechanisms, and the failure to recognise victims. These can be enacted quickly, using existing legislative processes, and they will create immediate protection forcitizens while longer-term constitutional and treaty work proceeds.

The legislative reform proposals for Belgium demonstrate what national legislation can achieve in the near term. They create criminal offences for directed-energy assault, expand stalking provisions to cover organised covert harassment, reform intelligence oversight, establish parliamentary inquiry mechanisms, and create victim support pathways. These measures do not require constitutional amendment. They require political will.

At the regional level, the EU can move faster than any single member state. A Regulation on covert directed-energy interference, combined with amendments to the AI Act and a Directive on victim protection, would create a uniform framework across 27 countries. The EU has the legislative machinery, the technical expertise, and the democratic legitimacy to act. What it lacks is the recognition that the problem exists.

At the international level, the UNESCO Recommendation provides a foundation. The next step is a UN General Assembly resolution calling for negotiations on a binding treaty. The precedent of the Biological Weapons Convention—negotiated in nine years during the Cold War—shows that such treaties are possible when the threat is recognised as existential.

The choice of legislative level is not an either/or question. It is a question of sequencing, coordination, and political strategy. The goal is not to find the single correct level but to build a system of protection that operates at every level, leaving no gap through which the human mind can be violated without consequence.

CHAPTER 67

Model Legislation — Drafting the Neuro-rights Statute

The gap between principle and practice in neuro-rights is largely a gap in legislative drafting. The principles — mental privacy, cognitive liberty, protection from neurodiscrimination — are broadly accepted. The legislative text that would operationalise these principles in enforceable statutory terms is largely absent from national and international law. This chapter offers a model approach.

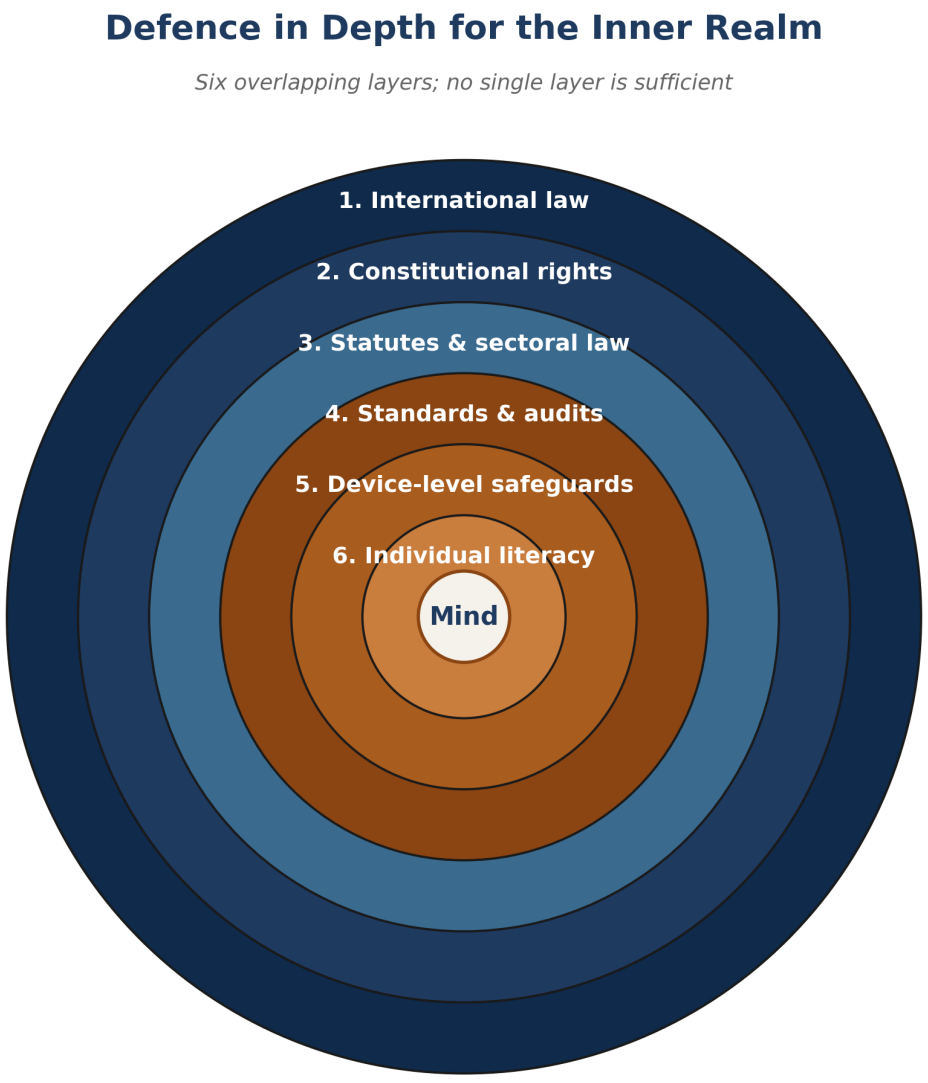


Figure 48.1. *Defence in depth for the inner realm. Six overlapping layers from international law to individual literacy; gaps in any one layer must be compensated by the others. The model statute drafted in this chapter sits in layer three.*

The Drafting Challenge

Legislative drafting for neuro-rights faces several distinctive challenges. The technology is rapidly evolving; statutory definitions must be future-proof. The harms are often diffuse and difficult to attribute; enforcement mechanisms must accommodate evidentiary difficulty. The actors range from consumer device manufacturers to intelligence agencies; the statute must apply uniformly across contexts. The international dimension is unavoidable; national legislation must interoperate with cross-border data flows and international standards.

The drafting approach proposed here addresses these challenges through five structural choices: technology-neutral definitions, layered protections proportional to harm, parallel public and private enforcement mechanisms, strong default protections with narrow exceptions, and clear interface with international human rights instruments.

Definition Layer

The foundation of any neuro-rights statute is the definition section. Three definitions are essential:

Neural data must be defined to include both raw neural signals (electrical, magnetic, optical, chemical, or other physical measurements of the central or peripheral nervous system) and any inferences derived from such signals through algorithmic or other processing. The inference component is critical: without it, regulated entities can launder raw neural data into "derived data" not covered by the statute.

Neurotechnology must be defined to include any device, system, or method capable of measuring, recording, or interpreting neural activity, or capable of stimulating, modulating, or directly altering neural activity. The definition must be technology-neutral, applying to invasive and non-invasive devices, electrical and other modalities, and consumer and medical applications alike.

Cognitive liberty must be defined as the right of every person to autonomy over their own mental processes, including the right to refuse neural surveillance, the right to refuse neural modification, the right to mental privacy, and the right to develop and express their thoughts without external interference.

The Inference Trap

A statute that protects only "neural data" without protecting algorithmic inferences contains a fatal loophole. Regulated entities can collect raw neural signals, immediately process them into inferred mental states (attention, emotion, intention), discard the raw data, and claim that the inferences are not "neural data" but proprietary analysis exempt from the statute. California's 2024 amendment explicitly addresses this by including inferences; Colorado's amendment does not. The difference determines whether the statute is enforceable in practice.

Substantive Prohibitions

The substantive prohibitions of the statute should be layered:

At the absolute layer: prohibition of non-consensual neural surveillance or modification, with no exceptions even in emergency or security contexts; prohibition of neuroweapons (devices designed primarily to cause harm through direct neural effect); prohibition of the use of neural data in criminal proceedings against the data subject; and prohibition of neurodiscrimination in employment, education, housing, insurance, and access to public services.

At the consent-based layer: collection of neural data subject to specific, informed, freely given consent, with consent invalid in coercive contexts such as employment and education without specific protections; processing of neural data limited to purposes specified at the time of consent; retention of neural data limited to the period strictly necessary for the consented purpose.

At the procedural layer: mandatory data protection impact assessment for any deployment of neurotechnology; mandatory transparency reporting on neural data collection, processing, and disclosure; mandatory independent technical auditing for high-risk deployments; mandatory tamper-evident logging for all neurotechnology capable of recording or stimulating neural activity.

Criminal Offences

The statute should establish specific criminal offences distinct from existing categories:

Non-consensual neurotechnological interference — the use of any device or method to monitor, record, stimulate, or modulate the neural activity of another person without that person's informed consent. Penalty proportional to other invasions of bodily integrity; aggravated penalty where the actor is a state agent, where the victim is a child or vulnerable adult, or where the interference produces lasting harm.

Directed-energy assault — the deliberate use of electromagnetic, acoustic, optical, or other directed-energy devices to cause physical pain, psychological distress, sleep deprivation, or cognitive interference. The offence should apply whether or not the use leaves externally visible injuries, and should not require proof that the perpetrator intended specific symptoms.

Coordinated neurotechnological harassment — the planning, organising, or participation in a coordinated campaign of neurotechnological interference involving multiple actors. The offence should apply to networks and command structures, not only to individual actors at the point of contact.

Trafficking in neural data — the unauthorised collection, sale, transfer, or use of neural data, with aggravated penalty where the data has been combined with other identifying information or used to make decisions affecting the data subject.

Civil Remedies

Criminal prosecution will not address every violation. The statute must provide civil remedies that data subjects can invoke directly:

A private right of action for any violation of the statute, with statutory damages set high enough to support enforcement (a model is the $1,000-$5,000 per violation framework of the Illinois Biometric Information Privacy Act). Injunctive relief should be available to prevent ongoing violations. Class action procedures should accommodate the typically dispersed nature of violations.

For state-actor violations, a specific civil rights cause of action should be available, with damages, injunctive relief, and attorneys' fees. Sovereign immunity defences should be narrowed or waived for violations of neuro-rights, consistent with the absolute nature of the rights at stake.

Institutional Provisions

The statute should establish or designate institutional infrastructure for enforcement:

A specialised neuro-rights regulatory body with technical capacity (RF spectrum analysis, neuroimaging, forensic engineering), investigative authority, and rule-making power. The body should be structurally independent of intelligence services and military commands that are themselves potential subjects of investigation.

A complaints procedure with statutory deadlines, technical investigation capability, and judicial review of dispositions. Complaints involving state actors should be subject to enhanced procedures including reference to parliamentary oversight bodies.

Victim support services with cause-neutral validation, trauma therapy access, and practical advocacy. The services should be available to complainants regardless of whether their specific claims are ultimately substantiated.

International Interface

National neuro-rights legislation must interface with the international framework:

Recognition of neural data as a category of personal data protected under cross-border data transfer rules, with restrictions on transfer to jurisdictions lacking equivalent protection. Recognition of foreign judgments enforcing neuro-rights, with appropriate public-policy reservations. Mutual legal assistance for cross-border investigations of neuro-rights violations.

The statute should explicitly incorporate the standards of the UN Human Rights Council Advisory Committee report A/HRC/57/61, the UNESCO Recommendation on the Ethics of Neurotechnology, and any subsequent international instruments. Where international standards exceed national provisions, the international standard should apply.

The Drafting Roadmap

A coordinated drafting effort should produce model legislation at several levels: national statute templates for adaptation to common-law and civil-law jurisdictions; supranational regulation models for the EU and other regional bodies; international convention text for adoption by the UN General Assembly; and constitutional amendment text for incorporation into national constitutions.

Such an effort is currently underway through a consortium of human rights organisations, academic institutions, and affected populations. The output should be made available as a legislative resource for any government considering neuro-rights legislation, with clear documentation of the policy choices embedded in each provision and the alternatives considered.

CHAPTER 67-bis

The Ideal Legislative Stack — Complete Model Statutes at Every Level

Chapters 65 to 68 introduced the legislative agenda, the principal levels at which it must be addressed, and the architecture of an international treaty. This chapter consolidates and extends. It offers a complete, level-by-level set of model legislative texts — drafted at the level of provisions rather than bills, but with enough specificity to be useful to a legislator, a civil servant or a civil-society advocate building a national or regional campaign.
We address all six tiers of the legislative stack identified in Chapter 18-bis and Chapter 66: supra-national treaty, regional framework, national constitution, national statute, sectoral regulation, and sub-national. At each tier we provide the substantive provisions we judge essential, drafted as legislative language. The drafting is intentionally compact; specific national adaptation will require longer prose, and translations into the languages of jurisdictions adopting the provisions will require care that no English-language draft can supply.

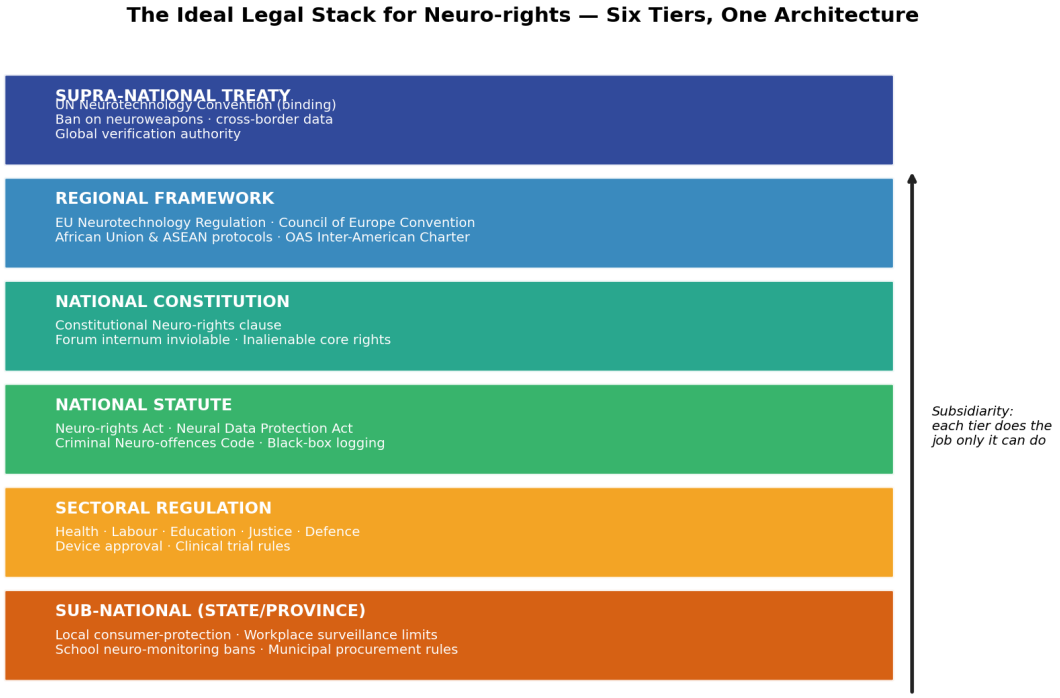


Figure 67bis.1. The ideal legal stack — six tiers, one architecture. Each tier does the job only it can do; gaps in any tier must be compensated by the others. The model provisions in this chapter aim to populate each tier with the minimum substantive content required for the stack to work as a whole.

Tier 1 — Supra-national: The Model UN Neurotechnology Convention

The supra-national tier requires a binding international convention. We propose its title as the United Nations Convention on the Protection of the Human Mind from Neurotechnological Harm (in short, the Neurotechnology Convention). It should be modelled on the Chemical Weapons Convention in structure, with substantive prohibitions, a verification regime, an organisation, and a dispute-settlement mechanism.

Preamble

The States Parties to this Convention, considering that recent developments in neurotechnology have created the technical possibility of reading and writing to the human mind in ways not previously possible; recognising that the human mind has traditionally and rightly been understood as the inmost sphere of personal sovereignty; affirming that the Universal Declaration of Human Rights and the International Covenants on Civil and Political Rights and on Economic, Social and Cultural Rights provide the foundational protections that this Convention seeks to extend; determined to prevent the development, production, stockpiling, transfer or use of neurotechnologies designed or adapted to covertly impair, control or manipulate cognition; and committed to the equitable global availability of neurotechnology's beneficial applications under conditions of human dignity, have agreed as follows.

Substantive Articles (Selected)

Article 1 (Scope). This Convention shall apply to all neurotechnologies, defined as devices, systems, substances or techniques that measure, interpret, modulate or stimulate the central or peripheral nervous system of a natural person, including the algorithmic derivation of inferences from the resulting data.

Article 2 (General Obligations). Each State Party undertakes never to develop, produce, otherwise acquire, stockpile, retain, transfer, or use any neurotechnology designed or adapted to: (a) measure neural or mental activity of a person without their free, specific, informed and revocable consent; (b) modulate or stimulate the nervous system of a person without such consent; (c) impair, suppress or alter the cognitive autonomy of a person.

Article 3 (Inalienable Core). The following protections are inalienable and may not be waived by consent, contract, terms of service or any other instrument: (a) the right to be free from involuntary neural reading or writing; (b) the right of children below the age of majority to be free from neurotechnological monitoring of attention, emotion or cognition in educational settings; (c) the right of workers to be free from neurotechnological monitoring as a condition of employment; (d) the right of persons deprived of liberty to be free from neurotechnological interrogation, monitoring or modification; (e) the right to be free from neurotechnological influence intended to alter electoral choice or political belief.

Article 4 (Neural Data Protection). Each State Party shall ensure in its national law that neural data — including both raw signals and algorithmic inferences derived from neural data — is treated as a special category of personal data requiring the highest standards of consent, security, purpose limitation and accountability. Neural data shall not be subject to commercial transfer, sale or licensing except as permitted by the laws of the State Party in conformity with this Convention.

Article 5 (Black-Box Logging). Each State Party shall require that all neurotechnologies placed on the market or used within its jurisdiction maintain tamper-evident logs of every neural measurement and stimulation event, of the parameters used, and of the destinations to which the resulting data is transmitted. Such logs shall be accessible to the subject of the neurotechnology, to designated regulators, and to independent forensic experts in the event of a dispute.

Article 6 (Verification Regime). The Organisation for the Prohibition of Neurotechnological Harm (OPNH), established under this Convention, shall maintain a system of declarations, routine inspections, and challenge inspections analogous to those established by the Chemical Weapons Convention. State Parties shall declare all neurotechnology development programmes, manufacturing facilities, and significant deployments to the OPNH, and shall permit inspections in accordance with the verification protocol annexed to this Convention.

Article 7 (Cross-Border Data Flows). Cross-border transfers of neural data shall be permitted only between jurisdictions that have ratified this Convention and have implemented its substantive provisions, except where the transfer is for the personal use of a natural person within their own control.

Article 8 (Individual Criminal Liability). Each State Party shall ensure that responsible individuals — including senior officials and corporate officers — face individual criminal liability for the development, deployment or use of prohibited neurotechnology. Such

offences shall be subject to universal jurisdiction analogous to that established for torture and crimes against humanity.
Article 9 (Victim Recognition and Redress). Each State Party shall establish in national law mechanisms for the recognition of victims of neurotechnological harm, for redress including compensation, and for the rehabilitation of affected persons. A Victim Trust Fund shall be established under the OPNH for cases in which national mechanisms are inadequate.
Article 10 (Whistleblower Protection). Each State Party shall provide effective protection in national law for persons who report violations of this Convention, including protection against employment retaliation, criminal liability for the act of reporting, and physical or psychological reprisal.
Article 11 (Synthetic Telepathy). Direct technological transfer of semantic content between the brains of natural persons (synthetic telepathy) shall be permitted only under the conditions set out in Annex IV, including symmetric consent, mandatory black-box logging, receiver override, and the prohibitions specified in Chapter 78-bis of the supporting commentary.
Article 12 (Augmentation and the Right to Remain Unaugmented). Each State Party shall ensure that cognitive augmentation is broadly and equitably available to persons within its jurisdiction who freely choose it, and that no person is required to undergo augmentation as a condition of employment, education, political participation or access to public services, save in narrow categories explicitly justified by safety considerations and approved by the OPNH.
Articles 13-25 cover specific cooperation, dispute settlement, amendment, ratification, entry into force, depositary functions and other standard treaty provisions.

Tier 2 — Regional Framework Provisions

Regional instruments — for the European Union, the African Union, ASEAN, the Organization of American States, the Council of Europe — should implement the Convention's substantive provisions within the regional legal frame and add region-specific provisions. We illustrate with model language for the EU.

EU: A Neurotechnology Regulation

Building on the AI Act, the GDPR and the Medical Device Regulation, the EU should adopt a dedicated Neurotechnology Regulation. Key provisions:
Article 1 (Subject Matter). This Regulation establishes harmonised rules for the placing on the Union market, the putting into service and the use of neurotechnologies within the Union, including provisions on neural data protection, on the conditions for synthetic telepathy services, and on the integration of the Union's obligations under the United Nations Neurotechnology Convention.
Article 5 (Prohibited Practices). The following are prohibited within the Union: (a) the placing on the market, putting into service or use of any neurotechnology intended to read neural activity without the free, specific, informed and revocable consent of the subject; (b) the placing on the market, putting into service or use of any neurotechnology intended to modulate the nervous system of a person beyond the specific therapeutic purposes for which the device has been approved, without such consent; (c) the use of any neurotechnology in workplace or educational settings for the monitoring of attention, emotion or cognition; (d) the use of any neurotechnology in connection with the political process for the purpose of influencing electoral choice or political belief.
Article 6 (High-Risk Neurotechnologies). All neurotechnologies are presumed to be high-risk AI systems under the AI Act and are subject to its requirements for high-risk systems, plus the additional requirements of this Regulation.
Article 8 (Black-Box Logging). All neurotechnologies placed on the Union market shall maintain tamper-evident logs in accordance with the standards adopted by the European Standardisation Bodies under the mandate of this Regulation.
Article 11 (Neural Data Protection). The processing of neural data shall be subject to Article 9 of the General Data Protection Regulation as a special category of personal data. The

Commission shall adopt by 2028 implementing acts specifying the conditions for lawful processing of neural inferences derived by AI systems.
Article 15 (Children). The use of neurotechnology in educational settings for the monitoring of children is prohibited. This prohibition applies to all primary and secondary educational settings within the Union, public and private. Limited exceptions apply for specifically-approved clinical uses with individual parental consent.
Article 18 (Workers). The use of neurotechnology in workplace settings for the monitoring of attention, emotion or cognition is prohibited. Employers may not require employees to participate in neurotechnology-based wellness programmes. The use of neurotechnology for safety-critical purposes (e.g., fatigue detection in commercial aviation) is permitted under specific exemptions issued by the European Aviation Safety Agency or other competent authority.
Article 21 (Cross-Border). Neural data may be transferred to third countries only on the basis of adequacy decisions adopted by the Commission, of binding corporate rules approved by the European Data Protection Board, or of specific derogations expressly authorised by this Regulation.
Article 25 (Sanctions). Infringements of this Regulation are subject to administrative fines of up to four per cent of the total worldwide annual turnover of the preceding financial year. Repeated or particularly serious infringements may result in the prohibition of the relevant economic operator from operating within the Union.

Tier 3 — National Constitutional Provisions

The constitutional tier requires a clause that places neuro-rights beyond the reach of ordinary political bargaining. We propose model language modelled on the Chilean amendment, with additions reflecting the lessons of the Brazilian and Spanish processes and the Convention's framework.

Model Constitutional Clause

Article [X] (Neuro-rights). 1. Every person has the right to the integrity of their mind. This right includes: (a) the right to the inviolability of mental content; (b) the right to mental privacy; (c) the right to cognitive autonomy; (d) the right to personal identity; (e) the right to be free from discrimination on the basis of mental or neural characteristics.
2. Neurotechnology that reads or writes to the human mind shall be subject to law. The State shall ensure that such neurotechnology is developed and used in ways consistent with this Article and with the international obligations of [the State].
3. The following are prohibited absolutely and may not be authorised by any law or contract: (a) the involuntary reading of neural activity; (b) the involuntary modulation of the nervous system; (c) neurotechnological monitoring of children in educational settings; (d) neurotechnological monitoring of workers as a condition of employment; (e) neurotechnological interrogation, monitoring or modification of persons deprived of liberty; (f) neurotechnological influence intended to alter electoral choice or political belief.
4. Neural data is the inalienable property of the person to whom it pertains. Neural data shall not be subject to commercial transfer, sale or licensing. Donation of neural data for research or medical purposes is permitted under conditions specified by law.
5. Any person whose rights under this Article have been violated shall have effective access to a remedy, including compensation, deletion of the data in question, and the protection of the violated right from future infringement.

Tier 4 — National Statutory Provisions: The Model Neuro-rights Act

Below the constitutional clause, national law should implement the constitutional provisions in detail. A model Neuro-rights Act might contain the following structure and substantive provisions.

Definitions

"Neurotechnology" means any device, system, substance or technique that measures, interprets, modulates or stimulates the central or peripheral nervous system of a natural person.

"Neural data" means: (a) any data obtained by direct or indirect measurement of the nervous system of a natural person, including electrical, magnetic, optical, biochemical or any other modality; and (b) any algorithmic inference derived from such data.

"Cognitive autonomy" means the capacity of a person to form, hold and act on thoughts, beliefs and decisions free of external coercion or technological manipulation that bypasses the conscious will of the person.

"Synthetic telepathy" means the direct technological transmission of semantic content between the brains of natural persons.

Substantive Prohibitions

Section 10. No person shall use, attempt to use or facilitate the use of neurotechnology to read neural activity of a natural person without the free, specific, informed and revocable consent of that person. Violation is punishable by [significant criminal penalty].

Section 11. No person shall use, attempt to use or facilitate the use of neurotechnology to modulate the nervous system of a natural person beyond the specific therapeutic purpose for which the device has been approved without such consent.

Section 12. No employer shall require, request or facilitate the use of neurotechnology by an employee for any purpose related to the employee's employment. The use of neurotechnology by an employee for the employee's own purposes shall not be the basis for any adverse employment action.

Section 13. No school, school district, educational institution or related entity shall use neurotechnology to monitor the attention, emotion or cognition of a child below the age of majority.

Section 14. No person shall use neurotechnology in connection with the political process for the purpose of influencing the electoral choice or political belief of any natural person.

Section 15. No person deprived of liberty shall be subjected to neurotechnology for any purpose related to their detention, prosecution or punishment.

Civil Remedies

Section 20. Any person whose rights under this Act have been violated shall have a private right of action against the violating person, including for: (a) actual damages; (b) statutory damages of not less than [specified amount] per violation; (c) attorney's fees and costs; (d) injunctive relief; (e) deletion of all neural data unlawfully obtained; (f) any other relief the court deems just.

Section 21. Class actions are available under this Act on a representative basis, with class certification subject to the ordinary rules of [the jurisdiction's] civil procedure.

Criminal Offences

Section 30. The intentional violation of Sections 10 through 15 is a felony punishable by [substantial imprisonment] and fines. Where the violation is committed by a corporate entity, individual officers who knew or should have known of the violation are subject to criminal liability under the Park doctrine.

Section 31. Where the violation involves children, persons deprived of liberty, or other particularly vulnerable populations, the penalties of Section 30 are doubled.

Section 32. Where the violation is committed in a pattern affecting one hundred or more persons, the offence is a class [X] felony with corresponding penalties.

Institutional Provisions

Section 40. There is established the National Neurotechnology Commission, an independent agency with the powers specified in Chapter 70-bis of the supporting commentary. The Commission shall: (a) license high-risk neurotechnologies for placement on the market; (b)

conduct compliance audits; (c) impose administrative penalties; (d) maintain the National Neurodata Register; (e) coordinate with international bodies including the OPNH.
Section 41. The Commission shall be funded through annual appropriation supplemented by fees collected from regulated entities. The Commission's budget shall be sufficient to maintain technical staff at levels comparable to those of the [national medical-device regulator].
Section 42. The Commission shall publish an annual report on neurotechnology use, compliance, and incidents within the jurisdiction.

Implementation Schedule

Section 50. This Act shall take effect on [date]. Existing neurotechnologies on the market on the effective date shall comply with the technical requirements of this Act within twenty-four months. The Commission shall be established within six months of the effective date.

Tier 5 — Sectoral Regulation: Model Provisions for Specific Sectors

Below the national statute, sectoral regulators must adopt detailed rules specific to their domains. We illustrate with model provisions for three principal sectors.

Health and Medical Devices

Subpart A. All neurotechnologies intended for medical use are subject to the [national medical-device regulator]'s pre-market authorisation. Devices that read or write to the central nervous system are presumed Class III (highest risk) and subject to the corresponding regulatory pathway.
Subpart B. Clinical trials of neurotechnology must comply with: (a) the consent provisions specific to neurotechnology trials as adopted by the [national institutional review board oversight body]; (b) the data-management provisions ensuring participant control of their neural data; (c) post-trial obligations including device explantation rights and ongoing care provisions.
Subpart C. The medical use of neurotechnology in emergency contexts where consent cannot be obtained is permitted subject to the substituted-judgement and best-interests standards specified by [the jurisdiction's emergency-care statutes], with documentation of the assessment and subsequent ratification by the patient or their decision-maker.

Labour and Workplace

Subpart A. The labour inspectorate shall conduct random compliance audits of workplace use of neurotechnology. Compliance audits shall include both technological inspection (to confirm absence of unauthorised devices) and employee interviews (to confirm absence of coercive practices).
Subpart B. Employees who report violations of the Neuro-rights Act in their workplace are protected against retaliation under [the whistleblower-protection statute] and are entitled to legal representation funded by the [labour department's] Whistleblower Support Fund.
Subpart C. The safety-critical exception to the workplace prohibition (Section 12 of the model Act) is available only on application to the labour inspectorate, with documentation of the safety necessity and with provisions ensuring proportionality and minimum-intrusiveness.

Education

Subpart A. Educational institutions are prohibited from acquiring or using any neurotechnology that measures the attention, emotion or cognition of a child. Inadvertent acquisition (e.g., neurotechnology features built into general consumer electronics) does not constitute a violation provided the institution disables or removes the feature within a reasonable time of becoming aware.
Subpart B. Neurotechnology used by a child for the child's own purposes outside the school setting (e.g., a personal sleep-tracking device) does not constitute a violation by the school.

However, the school shall not use, request access to, or facilitate the transfer of data from such devices.
Subpart C. The educational use of neurotechnology in tertiary education (above the age of majority) is subject to specific consent requirements adopted by the [national tertiary-education regulator].

Tier 6 — Sub-national and Municipal Provisions

At the sub-national level, states and provinces should adopt provisions that supplement the national framework with local-context-specific protections. Municipalities should adopt rules governing public procurement and public-space use of neurotechnology. The recent US state legislation (Colorado, California, Montana, Connecticut) provides templates for state-level action; the proposed Berkeley and Oakland municipal codes provide templates for municipal action.
The sub-national tier is also the level at which class actions are most effectively prosecuted in some jurisdictions, where representative-action mechanisms are most flexible, and where the political climate may permit more aggressive protection than the national consensus allows. The Chilean Emotiv case, the early Colorado neurodata cases, and the anticipated Brazilian state-level cases illustrate the model.

The Stack as a Whole

Each tier of the stack is necessary; none is sufficient. The supra-national Convention sets the prohibitions that no state can violate and supplies the international enforcement layer. The regional regulation harmonises within the bloc and operationalises the Convention. The constitutional clause entrenches the protection at the highest level of domestic law. The national statute supplies the operative rules and remedies. The sectoral regulation translates the rules into product specifications. The sub-national tier captures the local-context variations and political opportunities.
Gaps in any one tier must be compensated by the others. A jurisdiction without a constitutional clause can achieve substantial protection through aggressive sectoral regulation, as the EU has shown for data protection. A jurisdiction with a constitutional clause but weak statutory implementation will see the constitutional protection eroded by uncertainty. The Chilean Emotiv case worked because the constitutional clause and the data-protection statute together delivered a remedy; either alone might have failed.
The implementation of this stack — by whom, with what timelines, and at what cost — is the subject of Chapter 70-bis.

CHAPTER 68

The Architecture of an International Neurotechnology Treaty

Some neuro-rights violations cannot be effectively addressed by national legislation alone. Cross-border data flows, the deployment of neuroweapons across national boundaries, and the operation of multinational corporations all require international legal infrastructure. This chapter examines the architecture of an international neurotechnology treaty.

The Case for an International Instrument

Three categories of neuro-rights violation are inherently transnational. First, neural data collected in one jurisdiction routinely flows to processing centres in another and to data buyers in a third. Second, neuroweapons may be deployed from one jurisdiction against

targets in another, with the actors and the harm separated by international borders. Third, multinational neurotechnology companies operate across jurisdictions with different protections, creating opportunities for regulatory arbitrage.

Each of these requires international coordination. Even the strongest national legislation cannot effectively address violations that span borders without complementary international instruments. The current absence of any binding international neurotechnology treaty leaves a structural gap that no national action can fill.

Models from Existing International Law

The architecture of an international neurotechnology treaty can draw on three established models, each addressing a different dimension of the challenge:

The **Chemical Weapons Convention model** provides the template for prohibiting weapons of specific categories. The CWC establishes an absolute prohibition on chemical weapons, an inspections regime (the OPCW) with intrusive verification authority, and a destruction obligation for declared stockpiles. Adapted to neuroweapons, the model would prohibit specific categories of directed-energy and neural-interface weapons, establish verification authority, and require destruction of identified systems.

The **Council of Europe Convention 108 model** provides the template for data protection. Convention 108, the modernised version of which entered into force in 2018, establishes a framework for cross-border data protection cooperation that has been adopted beyond the Council of Europe membership. Adapted to neural data, the model would establish equivalent protection requirements, cross-border transfer restrictions, and a coordinating body.

The **Convention on the Rights of Persons with Disabilities (CRPD) model** provides the template for codifying emerging human rights. The CRPD synthesised existing rights principles into a binding instrument addressing a specific population. Adapted to neuro-rights, an analogous instrument would codify the five core neuro-rights, establish state obligations for implementation, and create a monitoring committee.

A Three-Pillar Treaty Structure

An effective international neurotechnology treaty should integrate these three models in a three-pillar structure:

Pillar 1: Prohibitions. The first pillar establishes absolute prohibitions on specific categories of neurotechnology deployment. The category list should include: neuroweapons designed to cause harm through direct neural effect; non-consensual neural surveillance of civilian populations; use of neural data in criminal proceedings against the data subject; and neural enhancement of combatants in armed conflict beyond defined safety thresholds.

Pillar 2: Data Protection. The second pillar establishes minimum protections for neural data, including consent requirements, purpose limitation, retention limits, and rights of access, correction, and deletion. The pillar should establish cross-border transfer restrictions and equivalence-based adequacy determinations for transfers to non-party states.

Pillar 3: Affirmative Rights. The third pillar codifies the five core neuro-rights as binding state obligations: mental privacy, personal identity, free will, fair access to cognitive enhancement, and protection from algorithmic bias. State obligations should include legislative implementation, institutional infrastructure, victim support, and access to remedy.

Verification and Enforcement

The verification challenge for an international neurotechnology treaty is substantial. Unlike chemical weapons, neuroweapons may leave no physical residue. Unlike traditional surveillance, neural data flows are often invisible. Unlike biological pathogens, neural interface technologies have ambiguous dual-use status.

A workable verification regime requires several mechanisms operating in parallel:

48. A declaration regime requiring states parties to declare their neurotechnology research and development programmes, including dual-use programmes with civilian applications.
49. A technical inspections capacity, operating through an international body, with authority to investigate alleged violations through forensic engineering, RF spectrum analysis, and review of technical documentation.
50. A complaints mechanism through which individuals can bring allegations of treaty violations directly to the international body, with appropriate procedural protections for states' national security interests.
51. A reporting and review mechanism modeled on the universal periodic review of the Human Rights Council, with regular peer assessment of compliance.
52. Sanctions for non-compliance, including suspension of treaty benefits, technical assistance, and — in serious cases — referral to the Security Council.

The Political Pathway

No international treaty arises spontaneously. The political pathway to a neurotechnology treaty requires several elements:

A coalition of states willing to lead. Historical precedents suggest that effective treaties often emerge from coalitions of middle powers willing to lead despite opposition from major powers. The Ottawa Process for landmines, the Cluster Munitions Convention, and the Treaty on the Prohibition of Nuclear Weapons followed this pattern. A comparable coalition for neuro-rights might centre on the states that have been most active to date: Chile, Spain, the European Union members, and a selection of middle powers from Africa, Asia, and Latin America.

Civil society mobilization. The treaties referenced above each emerged from sustained civil society campaigns. A neuro-rights treaty will require equivalent mobilization: documentation of harms, public advocacy, expert testimony, and political pressure across multiple states simultaneously.

Scientific and technical legitimacy. The treaty drafting process must integrate scientific and technical expertise. The OPCW maintains a Scientific Advisory Board; the Convention on Certain Conventional Weapons engages with expert communities; a neurotechnology treaty requires comparable infrastructure.

A drafting venue with credibility. The choice of venue — the UN General Assembly, the Conference on Disarmament, the Human Rights Council, or a stand-alone diplomatic conference — has substantive consequences. Each venue brings different participants, procedures, and political dynamics. The choice should be made strategically based on the political conditions of the moment.

The Timeline

International treaty processes typically span five to fifteen years from initiation to entry into force. The Chemical Weapons Convention took two decades from initial proposal to entry into force; the Treaty on the Prohibition of Nuclear Weapons took roughly a decade; the Convention on Certain Conventional Weapons protocols have taken varying periods. A neurotechnology treaty should be assumed to require a similar timescale.

This timescale poses a critical challenge: the technology is evolving faster than the treaty process. By the time a treaty enters into force, the technical landscape may have shifted substantially. The treaty architecture must therefore include amendment mechanisms, technical annexes that can be updated through procedures less cumbersome than treaty amendment, and explicit incorporation of evolving scientific knowledge.

The alternative — waiting for the technology to mature before beginning the treaty process — would guarantee that the treaty arrives too late. The lesson from previous arms-control efforts is that the political opening for prohibition typically appears before the technology has been fully normalised. Once normalised, prohibition becomes much more difficult. The window for an international neurotechnology treaty exists now; it may close within a decade.

CHAPTER 69

Immediate Action Required — A Call to Five Sectors

The legislative gap analysis is complete. The missing framework is mapped. The question that remains is not *what* must be done, but *who* must do it. This chapter addresses five sectors directly, specifying the actions each must take in the immediate term — the next twelve to eighteen months — to close the critical gaps before neurotechnology advances beyond the reach of law.

Law Enforcement

Police and investigative agencies are the first line of defense against neurotechnology abuse, yet they are currently the least prepared. Officers who receive complaints of directed-energy harassment, neural surveillance, or organized stalking have no standard protocol, no training, and no technical support. The result is systematic dismissal: victims are turned away, their reports unrecorded, their experiences attributed to mental illness.

Immediate actions required:

First, every police zone must establish a clear intake protocol for complaints involving neurotechnology or directed energy. The protocol must require that the complaint be logged and assigned a case number regardless of the officer's personal assessment of its plausibility. The existing stalking provisions of the Belgian Criminal Code (Articles 442bis–442ter) provide a legal basis for recording such complaints; what is missing is the operational instruction to use them.

Second, law enforcement agencies must develop partnerships with technical experts capable of conducting preliminary assessments. This means identifying engineers or physicists — whether within the force, at universities, or through contract — who can perform basic RF spectrum analysis at a complainant's residence, document anomalous electromagnetic readings, and provide expert reports admissible in court.

Third, police leadership must issue guidance recognizing that coordinated multi-actor harassment, including technology-mediated harassment, falls within the scope of existing stalking offences. The legislative reform proposals developed by TI One Voice explicitly recommend that "technology-mediated harassment including GPS tracking, drone surveillance, acoustic harassment devices, and electronic sabotage of the victim's devices or infrastructure" be included in the definition of *belaging/harcèlement*. Until such legislative amendments pass, police can and should interpret existing provisions to cover these modalities.

Fourth, specialized units must be created or designated within national police forces to handle neurotechnology-related complaints. These units require security clearances appropriate to the subject matter, access to forensic equipment, and authority to coordinate across jurisdictions when the technology in question operates at a distance.

Executive Branch

The executive branch — presidents, prime ministers, cabinets, and ministries — holds the power to initiate legislative action, allocate resources, and set the tone for public discourse. Inaction at this level is the single greatest obstacle to neuro-rights protection.

Immediate actions required:

First, each national government should commission a cross-departmental review of existing legal protections against neurotechnology abuse. The review should involve the ministries of justice, interior, health, defense, and science, and should produce a public report within six months identifying gaps and recommending legislative action. Belgium's existing legal architecture — the Constitution's Article 22 (right to privacy), the Criminal Code, the Organic Law on Intelligence Services — provides a baseline; the review must assess whether these instruments cover directed-energy assault and neuro-technological interference with cognition.

Second, governments should establish or designate a technical investigative capacity at the national level, competent in radio-frequency and electromagnetic forensics, signal analysis, and neurotechnological assessment. This unit must be independent of the intelligence services and accessible to citizens who report technological targeting. The model exists in the legislative reform proposals: a specialised diagnostic centre "equipped with RF spectrum analysis, audiological assessment, and neuroimaging capabilities, to provide independent medical-technical evaluation of TI claims."

Third, the executive must mandate annual public reporting by intelligence oversight bodies on the acquisition, possession, and deployment of directed-energy and neurotechnology capabilities by state security and military intelligence agencies. Transparency is the precondition for accountability.

Fourth, governments should endorse and actively support the UNESCO Recommendation on the Ethics of Neurotechnology, adopted by more than 190 member states in November 2025, and commit to implementing its principles in national law. This is soft law, but it carries moral and political weight; a government that endorses the Recommendation and then fails to act on it can be held to account by its own citizens.

Legislative Branch

Parliaments hold the power to create law. The five critical gaps identified in Chapter 47 — no definition of neuro-crimes, no neurodata protection, no neuroweapons ban, no oversight mechanism, no victim recognition — can only be closed through legislation.

Immediate actions required:

First, introduce and pass a criminal offence covering the deliberate use of electromagnetic, acoustic, or other directed-energy devices to cause physical harm, psychological distress, sleep deprivation, or cognitive interference. The proposed Article 422quater for the Belgian Criminal Code provides a model: a definition of "directed-energy device" broad enough to

encompass microwave, millimetre-wave, infrasonic, ultrasonic, and laser-based systems; aggravating circumstances for public officials and sustained campaigns; and penalties ranging from one to ten years imprisonment.

Second, amend data protection laws to define "neural data" to include not only raw brain signals but also the algorithmic inferences drawn from them. This is the single highest-value legislative fix, as the source material makes clear: "US states have already split on exactly this point, which makes the definition the most consequential clause in any bill."

Third, establish a parliamentary inquiry commission to investigate the scale and nature of reported covert harassment using neurotechnology and directed energy. The model is the Belgian Dutroux Commission (1996) and Rwanda Commission (1997): a commission with subpoena power, technical expertise, and a mandate to produce a public report with binding legislative recommendations within 18 months.

Fourth, mandate that neural data and mental integrity are expressly protected within existing data-protection and artificial-intelligence frameworks — the GDPR and the AI Act at EU level, and equivalent national instruments. The AI Act's provisions on cognitive manipulation (Article 5(1)(a)) do not extend to hardware-based neuro-interference; this gap must be closed.

Fifth, create a citizen complaint pathway within intelligence oversight bodies specifically for allegations of technology-facilitated harassment, with statutory response deadlines — 90 days for acknowledgement, 12 months for preliminary findings.

Judiciary

Courts will ultimately adjudicate the boundaries of neuro-rights. But judges cannot rule on what they do not understand. The judiciary faces a steep learning curve.

Immediate actions required:

First, judicial training institutes must develop programs on neurotechnology and directed energy for judges at all levels. The curriculum should cover the basic science (what RF energy can and cannot do to human tissue), the existing legal framework (the Charter of Fundamental Rights, the ECHR, the GDPR, the AI Act), and the evidentiary challenges (how to assess expert testimony on bioeffects, how to evaluate claims of remote neural interference).

Second, courts must develop protocols for handling cases involving neurotechnology evidence. This includes standards for the admissibility of expert testimony on bioeffects, guidance on the weight to be given to victim testimony when the harm is invisible to conventional forensics, and procedures for protecting victims from further harassment during legal proceedings.

Third, the judiciary should issue practice directions recognizing that the absence of externally visible injuries does not preclude a finding of assault or torture where the mechanism of harm is directed energy. The existing jurisprudence under Article 3 of the ECHR (prohibition of torture) and Article 4 of the Charter of Fundamental Rights provides a

foundation; what is needed is explicit recognition that these provisions apply to neurotechnological interference.

Fourth, courts should consider the evidentiary burden issue raised in the legislative reform proposals: once a victim can demonstrate a pattern of symptoms consistent with published bioeffects literature, the burden should shift to the alleged perpetrator to demonstrate that the technology in question was not used. This is not a reversal of the presumption of innocence; it is a recognition that the perpetrator has unique access to the evidence of their own actions.

Medical Sector

The medical sector is where victims of neurotechnology abuse most often seek help — and where they are most often failed. Healthcare professionals, lacking training on directed-energy bioeffects and neurotechnology-related symptoms, routinely misdiagnose technological torture as mental illness. The source material is explicit: "Medical professionals misdiagnose technological torture as mental illness." This must stop.

Immediate actions required:

First, medical associations and regulatory bodies must issue clinical guidance on the symptom profiles associated with directed-energy exposure and neurotechnological interference. These include tinnitus, unexplained burns, sleep disruption, cognitive fog, and auditory phenomena consistent with synthetic telepathy (voice-to-skull transmission). The guidance must require clinicians to consider environmental and technological causes before issuing psychiatric diagnoses for these symptom profiles.

Second, hospitals and clinics must establish referral pathways to specialised diagnostic centres equipped with RF spectrum analysis, audiological assessment, and neuroimaging capabilities. These centres should provide independent medical-technical evaluation for individuals who report technological targeting, and should produce reports that can be used in legal proceedings.

Third, medical schools and continuing education programs must incorporate neurotechnology and directed-energy bioeffects into their curricula. This is not fringe science; the bioeffects literature is extensive, peer-reviewed, and published in reputable journals. The IEEE C95.1 standard and ICNIRP guidelines provide reference points for non-thermal effects. Medical professionals should know this literature exists and how to apply it.

Fourth, the medical sector must formally recognize "Targeted Individual" status within patient records and diagnostic categories. This does not require accepting any contested factual claim about the cause of the symptoms; it requires recognizing that a growing number of patients present with a consistent cluster of experiences that they attribute to technological means, and that these patients deserve competent, respectful care regardless of the ultimate etiology.

Fifth, healthcare institutions must implement protocols to protect patients who report technological targeting from discrimination, involuntary psychiatric commitment, and other

forms of institutional harm. The source material documents that victims are “systematically gaslighted by institutions”; this pattern must be identified and eliminated through clear policies and accountability mechanisms.

These five sectors — law enforcement, executive branch, legislative branch, judiciary, and medical sector — hold the keys to closing the neuro-rights protection gap. Each sector has specific, achievable actions it can take in the next twelve to eighteen months. None of these actions requires constitutional amendment. None requires a binding international treaty. None requires technology that does not exist.

What they require is recognition: that the problem is real, that the victims are credible, and that the law must evolve to protect the last private place — the human mind.

PART XIII

Oversight and Governance of Neurotechnology

What form should oversight take? The question is deceptively simple. In practice, oversight of neurotechnology must operate at multiple levels simultaneously: technical audit of devices and emissions, ethical review of research and deployment, independent investigation of complaints, and democratic accountability for state actors who possess these capabilities. No single institution can fulfill all these functions. The challenge is to design a system of overlapping and mutually reinforcing mechanisms that together create meaningful accountability.

The source material provides concrete models. The Belgian legislative reform proposals call for a specialized technical assessment unit within the Comité R (Standing Intelligence Agencies Review Committee), staffed with engineers and physicists holding appropriate security clearances, capable of independently auditing electromagnetic emissions from suspected government installations. This is not a hypothetical recommendation; it is a practical response to the documented failure of existing oversight bodies to detect or investigate neurotechnological abuse.

At the European level, the proposals envision a new EU Regulation establishing an absolute prohibition on the use of directed-energy weapons and neuro-technological devices against any person within EU territory, with strictly defined law-enforcement exceptions subject to judicial authorization, proportionality review, and mandatory logging. The regulation would also mandate registration of all directed-energy devices above defined power thresholds with a new EU-level registry, managed by the European Defence Agency or a designated civilian authority.

The oversight architecture must include several essential components.

Independent Neurotechnology Commissions. At the national level, each country should establish a permanent commission with statutory independence, adequate funding, and technical expertise. The commission's mandate should include: maintaining a public register of authorized neurotechnology research and deployment; investigating complaints from citizens who allege neurotechnological harm; conducting proactive audits of government agencies that possess neurotechnology capabilities; and reporting annually to parliament with findings and recommendations. The commission must have subpoena power, access to classified information, and the authority to refer cases for prosecution.

Technical Audit Mechanisms. Oversight cannot function without the capacity to detect and verify. This requires: spectrum monitoring stations capable of detecting anomalous electromagnetic emissions across civilian frequency bands; forensic protocols for collecting and analyzing evidence of directed-energy exposure; accredited laboratories that can test devices against claimed specifications; and a whistleblower protection regime that enables engineers and scientists to report misuse without fear of retaliation. The source material emphasizes that existing legal frameworks fail because "investigative bodies lack technical expertise to verify neurotechnological harassment." Technical audit mechanisms close this gap.

Ethics Committees with Real Power. Existing research ethics committees typically review proposals before the research begins, but lack authority to monitor ongoing deployment or investigate complaints after the fact. Neurotechnology oversight requires ethics committees with continuing jurisdiction: the power to halt deployments, order independent testing, and require remediation when harms are discovered. These committees should include not only scientists and ethicists but also community representatives and, where appropriate, individuals with lived experience of neurotechnological targeting.

Citizen Complaint Pathways. Any oversight system that does not provide an accessible route for citizens to report abuse is structurally incomplete. The source material calls for a "citizen complaint pathway within the Comité R specifically for allegations of technology-facilitated harassment by intelligence services, with statutory response deadlines (90 days for acknowledgement, 12 months for preliminary findings)." Similar pathways should exist at every level: local police, national oversight bodies, and international human rights mechanisms. Complaints must be taken seriously, investigated competently, and resolved in a reasonable timeframe.

Parliamentary Inquiry Commissions. When existing mechanisms fail, parliaments must have the power to launch independent investigations. The Belgian model proposes a parliamentary inquiry commission "modelled on the precedent of the Dutroux Commission (1996) and the Rwanda Commission (1997)," tasked with investigating the scale and nature of reported covert harassment, auditing relevant capabilities of intelligence services, reviewing the adequacy of existing legal protections, and producing a public report with binding legislative recommendations within 18 months. Such commissions have the power to compel testimony, demand documents, and focus public attention on systemic failures.

International Coordination. Neurotechnology does not respect national borders. A device operated in one country can affect a person in another. Oversight mechanisms must therefore include international coordination: mutual recognition of complaints, cross-border investigative cooperation, shared technical standards, and ultimately a global registry of authorized neurotechnology capabilities. The European proposals for mutual recognition of victim status across Member States provide a model that could be extended globally.

The source material makes clear that existing oversight mechanisms are inadequate. The UN Human Rights Council Advisory Committee has noted that "there are no discrete human rights policies to limit the misuse and abuse of neurotechnology." This regulatory vacuum is not an accident; it is the predictable result of technologies that have outpaced the legal and institutional frameworks designed to govern them.

Closing the vacuum requires more than new laws. It requires new institutions with the independence, expertise, and authority to make those laws meaningful. The oversight architecture proposed here — independent commissions, technical audit mechanisms, empowered ethics committees, citizen complaint pathways, parliamentary inquiries, and international coordination — provides a blueprint for building those institutions.

The alternative is a world where neurotechnology develops in a human rights vacuum, where victims have nowhere to turn, and where the most intimate aspects of human consciousness become the province of unaccountable power. That is not a future we should accept.

CHAPTER 70

The Shape of Oversight — Independent Neurotechnology Commissions

The previous chapter established *why* oversight is necessary. This chapter addresses *what form that oversight should take*. The question is not whether neurotechnology should be governed — the technologies are too powerful, too intimate, and too easily weaponized to remain ungoverned. The question is what institutional architecture can credibly perform the task.

Existing oversight mechanisms have proven inadequate. As the UN Human Rights Council Advisory Committee has documented, "there are no discrete human rights policies to limit the misuse and abuse of neurotechnology." This is not an oversight gap that can be closed by a single law or a single agency. It requires a layered institutional architecture at the national, regional, and international levels.

The Core Principles of Neurotechnology Oversight

Any credible oversight body must be built on five principles drawn from the experience of existing oversight mechanisms for intelligence services, weapons development, and emerging technologies:

Independence. The oversight body must be structurally separate from the entities it oversees. It cannot report to the same ministry that funds neurotechnology research, nor can its members be appointed by the same officials who authorize its deployment. The Belgian model of the Comité R — the Standing Intelligence Agencies Review Committee — provides one template: an independent body whose members are appointed by parliament, not the executive, and which reports directly to the legislative branch.

Technical Competence. Oversight of neurotechnology requires the ability to audit electromagnetic emissions, analyze neural data processing algorithms, evaluate directed-energy device specifications, and understand the bioeffects literature. This is not work that can be done by generalist lawyers or policy analysts alone. The oversight body must include engineers, physicists, neuroscientists, and forensic specialists with appropriate security clearances and the mandate to conduct independent technical assessments.

Enforcement Power. An oversight body that can only issue recommendations is not oversight; it is advisory. The body must have the authority to compel testimony, demand documents, order inspections, issue cease-and-desist orders, refer cases for prosecution, and — in extreme cases — order the seizure or destruction of unauthorized devices. These powers must be backed by judicial oversight to prevent abuse, but they must exist in substance.

Citizen Access. Victims of neurotechnology abuse currently have no formal complaint pathway. The oversight body must establish a mechanism for individuals to file complaints, receive acknowledgment within statutory deadlines, obtain preliminary findings within a reasonable period, and access remedies if violations are confirmed. The Belgian reform

proposals for the Comité R include precisely such a citizen complaint pathway, with 90-day acknowledgment deadlines and 12-month preliminary findings requirements.

Transparency and Accountability. The oversight body must report publicly on its activities, findings, and recommendations, subject only to legitimate national security redactions that must themselves be subject to independent review. Secrecy is the enemy of accountability, and the history of covert neurotechnology programs — from MKULTRA to the present — demonstrates that what is hidden is what is abused.

National Oversight: The Independent Neurotechnology Commission

At the national level, each country should establish an Independent Neurotechnology Commission (INC) with the following composition and powers:

Composition. The commission should have 7-11 members appointed by the legislature, not the executive, for fixed staggered terms of 6-9 years. Membership should include: a judge or former judge with experience in technology law; a neuroscientist with expertise in neurotechnology; an engineer with expertise in electromagnetic systems and directed energy; a human rights lawyer; an ethicist specializing in bioethics or technology ethics; a representative of civil society organizations working on technology and human rights; and a person with lived experience of neurotechnology-related harm.

Powers. The commission should have the authority to: maintain a registry of all neurotechnology devices and systems in use within the jurisdiction; conduct unannounced inspections of facilities where neurotechnology is developed, stored, or deployed; compel the production of documents and testimony; issue binding orders to cease unauthorized activities; refer cases for criminal prosecution; recommend legislative changes to parliament; and publish annual public reports on the state of neurotechnology governance.

Funding. The commission must be independently funded through a dedicated line in the national budget, not through the ministry responsible for technology development or national security. The funding should be sufficient to maintain a full-time technical staff, including RF engineers, forensic analysts, and data scientists. A common vulnerability of oversight bodies is that they are starved of resources precisely so that they cannot perform their functions effectively.

Technical Audit Unit. Within each INC, a specialized Technical Audit Unit should be established, staffed with engineers and physicists capable of: conducting RF spectrum analysis to detect unauthorized directed-energy emissions; auditing neural data processing algorithms for bias and unauthorized data extraction; evaluating device specifications against manufacturer claims; and developing forensic methodologies for detecting neurotechnology use after the fact. The Belgian reform proposals for the Comité R include precisely such a specialized unit, and this model should be replicated globally.

Regional Oversight: The European Model

At the regional level, the European Union provides the most advanced existing framework for transboundary oversight. The EU's Charter of Fundamental Rights, the GDPR, and the AI Act provide legal anchors, but none of these instruments currently addresses hardware-based neuro-interference or directed-energy weapons used against civilians within EU territory.

The proposed EU Regulation on the Prohibition of Covert Directed-Energy and Neuro-Technological Interference Against Persons would establish: an absolute prohibition on the use of directed-energy weapons and neuro-technological devices against any person within EU territory, except under strictly defined law-enforcement scenarios subject to judicial authorization; a ban on the development, manufacture, sale, and possession of devices specifically designed for covert harassment; mandatory registration of all directed-energy devices above defined power thresholds with an EU-level registry; and the establishment of a right to electromagnetic environmental integrity as an extension of the Article 7 Charter right to private life.

The European Defence Agency or a designated civilian authority would maintain the device registry and coordinate cross-border investigations. Eurojust would provide coordination for cases involving multiple Member States. The European Ombudsman would provide an additional complaint pathway for individuals who cannot obtain redress at the national level.

International Oversight: The Global Neurotechnology Authority

At the international level, the gaps are even more profound. No existing treaty body has the mandate or capacity to oversee neurotechnology. The UN Human Rights Council can issue resolutions and commission studies — as it did with Resolution 51/3 in 2022 and the resulting 2024 report — but it cannot inspect facilities, compel testimony, or enforce compliance.

The logical model is a treaty-based International Neurotechnology Authority (INA), analogous to the Organisation for the Prohibition of Chemical Weapons (OPCW) established under the Chemical Weapons Convention. The INA would: maintain a global registry of authorized neurotechnology capabilities; conduct inspections of declared facilities; investigate allegations of unauthorized use; provide technical assistance to national oversight bodies; and coordinate international responses to cross-border incidents.

Negotiating such a treaty would take years, and the political obstacles are formidable. But the trajectory of neurotechnology development — from consumer BCIs to directed-energy weapons to AI-driven neural surveillance — suggests that the need for such an institution will become undeniable within the next decade. The question is whether we build it proactively or reactively, after the first mass-casualty neuroweapons incident or the first international crisis triggered by unauthorized neural surveillance.

The Citizen Complaint Pathway

A critical function of any oversight architecture is the citizen complaint pathway. Victims of neurotechnology abuse currently have no formal mechanism to report violations, no body with the technical expertise to investigate their claims, and no avenue for redress if violations are confirmed. This is not merely a procedural gap; it is a human rights crisis.

The proposed pathway includes: a dedicated complaint intake unit within each INC, accessible by phone, online portal, and in-person at designated locations; statutory deadlines for acknowledgment (90 days) and preliminary findings (12 months); the right to independent technical assessment of the alleged violation; the right to legal representation throughout the process; the right to appeal adverse findings to an independent tribunal; and the right to compensation if violations are confirmed.

For victims whose complaints involve state actors or classified technologies, the pathway must include secure procedures that protect both the victim's safety and legitimate national security concerns. This is not an impossible balance. Intelligence oversight bodies in democratic countries already manage classified complaint procedures; the same mechanisms can be adapted for neurotechnology.

The Cost of Inaction

The source material makes clear that the current regulatory vacuum is not benign. It actively enables abuse. Without independent oversight bodies, victims have nowhere to turn. Without technical audit capabilities, violations go undetected. Without enforcement powers, even identified violations go unpunished. Without citizen complaint pathways, the most vulnerable are silenced.

The cost of establishing independent neurotechnologycommissions is measurable in money, personnel, and political capital. The cost of not establishing them is measured in human lives, mental integrity, and the erosion of democratic accountability.

The estimates for a fully functional national INC range from €5-15 million annually for a country of moderate size, depending on the scope of its mandate and the technical sophistication of its audit unit. For the European-level registry and coordination mechanism, the costs would be higher — perhaps €50-100 million annually — but still a fraction of what governments spend on neurotechnology research and development. The United States alone has invested billions in DARPA's neural interfacing programs; the cost of oversight is less than one percent of the cost of development.

The political obstacles are more significant. Intelligence agencies resist oversight. Defense contractors resist transparency. Technology companies resist regulation. And the very nature of covert neurotechnology programs means that the most dangerous capabilities are the ones most fiercely protected from scrutiny. But these obstacles are not insurmountable. The Chemical Weapons Convention was negotiated during the Cold War, against the opposition of powerful military establishments. The Biological Weapons Convention was signed despite the objections of the defense industry. The Ottawa Treaty banning landmines was achieved through the concerted action of civil society, middle powers, and international organizations.

The same coalition can be built for neurotechnology oversight. The victims — the Targeted Individuals whose testimony is documented in the source material — provide the moral urgency. The scientific community provides the technical expertise. The human rights community provides the legal framework. And the growing recognition among governments that unregulated neurotechnology threatens everyone — not just the targeted few — provides the political opening.

The shape of oversight is clear. The question is whether we have the will to build it.

CHAPTER 70-bis

Implementation — Existing Agencies and the New Agencies Needed

The previous chapter set out the substantive provisions of the ideal legislative stack. This chapter addresses the implementation question: which agencies will enforce the rules, what mandate expansions are required for existing agencies, and what new agencies must be created. A law without an institutional implementer is a wish; the institutional architecture is the difference between substantive protection and decorative prose.

The architecture proposed below is intentionally pragmatic. We begin with existing agencies — the data-protection authorities, medical-device regulators, telecommunications authorities, labour inspectorates, equality bodies, consumer-protection agencies, and education ministries that already operate in most jurisdictions — and identify the specific mandate expansions each requires. We then identify six new agency types that must be created or designated to fill the gaps that no existing agency can credibly fill.

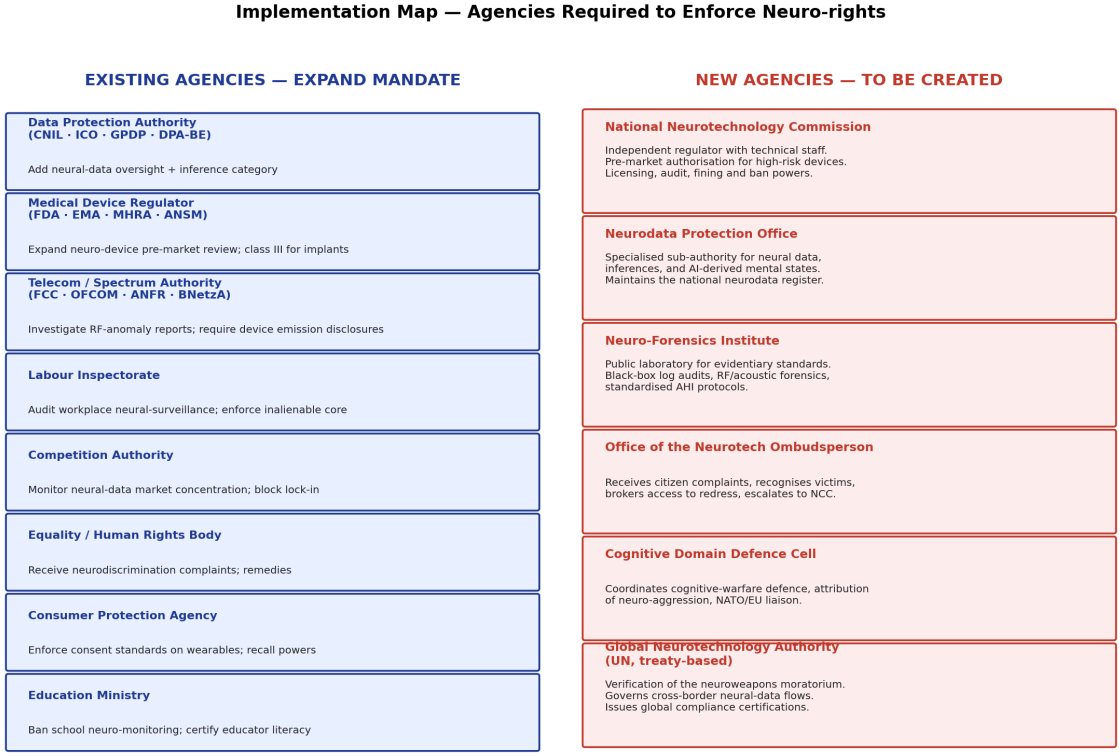


Figure 70bis.1. Implementation map for the ideal legislative stack. Eight existing agencies, with expanded mandates, supply most of the day-to-day enforcement; six new agencies fill gaps that no existing actor is well positioned to address.

Existing Agencies — Mandate Expansions

Data Protection Authorities

Data protection authorities (DPAs) — the CNIL in France, the ICO in the UK, the Garante in Italy, the AEPD in Spain, the BfDI in Germany, the FTC in the United States, the EDPB at EU level — are the existing agencies best positioned to take on most of the day-to-day enforcement of neural-data protection. Their existing mandate over personal data extends naturally to neural data; their existing technical capacity (such as it is) provides a foundation.

Required expansion: explicit jurisdiction over neural data, including algorithmic inferences derived from neural data; technical staff with neuroscience or neuroengineering training; the power to require pre-deployment data-protection impact assessments specifically for neural data; the power to order deletion of unlawfully obtained neural data (the Chilean precedent), with extraterritorial effect where the controller is subject to the DPA's jurisdiction; integration with the national medical-device regulator and the new National Neurotechnology Commission.

Medical Device Regulators

The FDA in the United States, the EMA and national agencies in the EU, the MHRA in the UK, the PMDA in Japan, the NMPA in China, the ANVISA in Brazil and the corresponding bodies in other jurisdictions regulate the medical-grade tier of neurotechnology. Their existing pre-market approval pathways, their post-market surveillance capacity, and their inspection authority provide much of the necessary infrastructure.
Required expansion: clearer jurisdiction over devices that straddle the medical and consumer tiers (this is the principal regulatory ambiguity); explicit pre-market review of black-box logging compliance; the power to require post-market reporting of long-term neural-data flows from manufacturer cloud services; integration with national data-protection authorities and the National Neurotechnology Commission.

Telecommunications and Spectrum Authorities

The FCC in the United States, OFCOM in the UK, the ANFR in France, the BNetzA in Germany, and corresponding bodies elsewhere regulate the radio-frequency spectrum and the certification of devices that emit on it. The relevance to neurotechnology is twofold: most consumer-grade neurotechnology emits on the spectrum and is subject to existing rules; alleged covert directed-energy devices, if they exist, also operate on the spectrum.
Required expansion: technical capacity to investigate reports of anomalous radio-frequency exposure (the Havana Syndrome-related capacity recommended in Chapter 40); explicit authority to inspect facilities suspected of housing prohibited neurotechnology; integration with the proposed Neuro-Forensics Institute and the international Organisation for the Prohibition of Neurotechnological Harm.

Labour Inspectorates

Labour inspectorates and equivalent bodies enforce workplace law. Their existing authority over workplace surveillance practices supplies most of the framework for enforcing the prohibitions on workplace neurotechnology specified in the model Act.
Required expansion: technical capacity to detect the presence of neurotechnology in workplace settings; jurisdiction over neurotechnology-based wellness programmes that constitute coerced workplace surveillance under the model Act; integration with whistleblower-protection bodies for receiving and acting on employee complaints.

Competition Authorities

Competition authorities — the European Commission's DG COMP, the FTC and DOJ in the United States, the CMA in the UK, the Bundeskartellamt in Germany, and corresponding bodies in other jurisdictions — should pay specific attention to the structural risks of concentration in the neurotechnology sector.
Required expansion: explicit guidance on merger review for neurotechnology firms (the recent Big Tech moves into BCI standardisation are a leading concern); investigation of vertical integration between device manufacturers, AI inference providers and downstream platforms; explicit treatment of neural data as a competition-relevant asset; coordination with national security and privacy regulators on dual-use technology.

Equality and Human Rights Bodies

National human rights commissions and equality bodies — the Equality and Human Rights Commission in the UK, the EEOC in the United States, the Défenseur des droits in France, the corresponding bodies in other jurisdictions — should explicitly take on neurotechnology-related discrimination claims.
Required expansion: the establishment of "neurodiscrimination" as a recognised category of discrimination, with explicit guidance on its application to employment, credit, insurance, education and access to public services; integration with the new Office of the Neurotech Ombudsperson for victim support.

Consumer Protection Agencies

Consumer protection agencies should enforce the consumer-tier provisions of the model Act, particularly the prohibitions on deceptive consent and on unauthorised data flows from consumer wearables.
Required expansion: dedicated consumer-tier neurotechnology team with technical capacity; explicit guidance on consent standards for consumer neurotechnology, including the prohibition on consent obtained through bundled terms of service; coordination with the Federal Trade Commission's equivalents on cross-border products.

Education Ministries

Education ministries should enforce the prohibition on neurotechnology in schools.
Required expansion: explicit guidance to school districts on permitted and prohibited uses; certification of educator literacy on neurotechnology issues; integration with consumer-protection agencies for cross-cutting deployments.

New Agencies — Required to Be Created

Six new agency types must be created at national level (and one at international level) to fill the gaps that no existing agency can credibly fill.

The National Neurotechnology Commission

The National Neurotechnology Commission (NNC) is the most important new institution. It is the central regulatory authority for neurotechnology within the jurisdiction, with the following characteristics:
Independence. The NNC must be independent of the executive in the same sense and to the same degree as the central bank, the data-protection authority and the medical-device regulator. The Commissioners must serve fixed terms, must be removable only for cause, must have professional qualifications in relevant fields, and must be appointed through a process that includes legislative and civil-society input.
Technical capacity. The NNC must employ a substantial technical staff with neuroscience, neuroengineering, AI, security and policy expertise. Staffing levels comparable to the national medical-device regulator are the minimum; staffing comparable to the data-protection authority is the target.
Powers. The NNC must have the power to: (a) license high-risk neurotechnologies for placement on the market; (b) deny or revoke licences; (c) conduct compliance audits with subpoena power; (d) impose administrative penalties up to a substantial fraction of global turnover; (e) refer matters for criminal prosecution; (f) issue binding standards and guidance; (g) maintain the National Neurodata Register; (h) coordinate with international bodies including the OPNH; (i) require pre-deployment ethical impact assessments; (j) commission independent red-team testing of products and AI models.
Funding. Sustainable, independent funding through a combination of annual appropriation, regulatory fees from licensed entities, and (in some jurisdictions) penalty income retained for technical capacity.
Accountability. The NNC must publish annual reports, must be subject to parliamentary or congressional oversight, must subject its own operations to data-protection and access-to-information requirements, and must operate a transparent complaint procedure.

The Neurodata Protection Office

Specialised sub-authority for neural data within the National Neurotechnology Commission (or in some jurisdictions integrated into the existing Data Protection Authority). Its specific responsibilities include: maintenance of the national neurodata register; oversight of cross-border neural-data flows; investigation of consumer complaints specific to neural data; technical guidance on the protection of inferences derived from neural data; coordination with international counterparts under the proposed Convention.

The Neuro-Forensics Institute

Public laboratory for the development and application of evidentiary standards in neurotechnology cases. The Institute's responsibilities include: maintenance of the standard protocols for evaluating black-box logs; the conduct of forensic analyses in disputed cases; the development of RF and acoustic-anomaly protocols (the Havana Syndrome-related capacity); the training of investigators in the National Neurotechnology Commission, in law enforcement, and in medical practice; the publication of standardised methodologies for use by independent practitioners.

The Institute is the closest analogue, in the neurotechnology domain, of the national forensic-science capacity that already exists for ballistic evidence, DNA analysis and forensic accounting. Its existence is necessary because no other institution combines the scientific rigour required with the independence from particular cases that evidentiary work demands.

The Office of the Neurotech Ombudsperson

An independent office receiving complaints from individuals who believe their neuro-rights have been violated, including those who self-identify as victims of covert or unauthorised neurotechnology. The Ombudsperson is the cause-neutral support institution recommended in Chapter 42, with the responsibilities of: receiving complaints and providing initial assessment; brokering access to the appropriate enforcement body (Data Protection Authority, NNC, criminal justice, etc.); facilitating access to clinical evaluation and (where indicated) medical treatment; supporting victims through the often-protracted process of remedy; publishing aggregate-level reports on patterns of complaint.

The Ombudsperson is the only institution in the architecture explicitly oriented toward the experience of the victim, regardless of the eventual determination of cause. It is the institutional embodiment of the cause-neutral, humane response described in Chapter 36.

The Cognitive Domain Defence Coordination Cell

A national security-related body coordinating defence against cognitive-warfare threats. Its responsibilities include: monitoring of disinformation, microtargeting and other cognitive-warfare techniques against the national population; coordination with the NATO Cognitive Warfare Coordination Centre and equivalent allied bodies; attribution of cognitive-warfare operations to responsible state and non-state actors; preparation of options for collective response under the proposed Convention.

In many jurisdictions, this cell would sit within an existing national security body (the National Security Council, the Ministry of Defence, the foreign intelligence service) rather than as a free-standing agency. The point is that the function be explicit and resourced, regardless of its location in the national security architecture.

The Global Neurotechnology Authority

At international level, the OPNH described in Chapter 67-bis's model Convention. The functions and powers of the OPNH should be modelled on the Organisation for the Prohibition of Chemical Weapons: declarations from State Parties, routine inspections, challenge inspections, dispute settlement, and technical assistance. The OPNH should be headquartered in a neutral jurisdiction; its inspectorate should be independent of any single state; its membership should be open to all UN Member States.

The OPNH's distinctive challenge is that the verification problem is harder than the chemical-weapons case. Many neurotechnologies have legitimate civilian uses; the dual-use boundary is more porous than in the chemical-weapons case; the technologies do not leave easily-detected physical traces in the way that chemical agents do. The OPNH's verification protocol will need to evolve over its first decade of operation, and the international community will need to accept that imperfect verification is better than no verification.

Coordination and Subsidiarity

The coordination problem across these agencies is the principal pragmatic challenge. The NNC should serve as the central coordinator at national level, with formal coordination mechanisms with each of the other relevant agencies. At international level, the OPNH coordinates with the OECD, UNESCO, WHO and other relevant bodies; bilateral and multilateral cooperation arrangements supplement the formal Convention machinery.
The principle of subsidiarity should govern the allocation of responsibilities. Routine consumer-tier enforcement belongs to the consumer-protection agency. Routine workplace-tier enforcement belongs to the labour inspectorate. Routine medical-tier enforcement belongs to the medical-device regulator. The NNC handles the matters that fall between or above these — high-risk product licensing, cross-tier disputes, novel regulatory questions. The OPNH handles the international-level matters and the disputes that cross jurisdictional boundaries.

Costs and Phasing

The total cost of establishing and operating this architecture is substantial but not enormous. The NNC's annual operating budget, in a major economy, is plausibly in the range of [200-500 million USD-equivalent], comparable to the national medical-device regulator. The Neuro-Forensics Institute is in the range of [50-100 million]. The Ombudsperson and other smaller offices are in the range of [10-30 million each]. The OPNH's annual budget, assessed across all State Parties, is plausibly comparable to the OPCW's annual budget (approximately 70 million Euro).
A realistic phasing for a national jurisdiction starting from scratch is: in year one, designate the existing agencies' expanded mandates and begin the legislative process for the NNC; in years two and three, establish the NNC and the Ombudsperson; in years three to five, establish the Neuro-Forensics Institute and complete the integration with existing agencies; in years five to ten, accumulate operational experience and refine the regulatory machinery. The OPNH, depending on the diplomatic trajectory, plausibly enters into force within ten to fifteen years of the present moment.

A Working Conclusion

The implementation architecture proposed above is large, complex, and expensive. It is, on our analysis, the minimum required for the substantive provisions of the model legislation to mean what they say. A jurisdiction that passes the model Neuro-rights Act but does not establish the NNC will have produced symbolic protection; the rights without the institutional implementer will not protect citizens. Conversely, a jurisdiction that establishes the NNC and the supporting machinery — even with somewhat weaker substantive law — will produce more protection in practice than a jurisdiction that does the opposite.
The historical analogy is the establishment of the data-protection regime in Europe through the 1990s and 2000s. The substantive law evolved over decades; the institutional infrastructure (the national DPAs, the European Data Protection Supervisor, the European Data Protection Board) evolved alongside. Without the institutional infrastructure, the substantive law would have been unenforceable; without the substantive law, the institutional infrastructure would have lacked purpose. The same is true for neuro-rights, and the next decade is when the matching of the two must be done.

CHAPTER 71

Technical Audit Mechanisms — How to Verify Compliance

Oversight without verification is theater. A neurotechnology commission may have the finest legal mandate, the most distinguished commissioners, and the most eloquent mission statement, but if it lacks the technical capacity to determine whether a directed-energy device is operating in a residential neighborhood, whether an intelligence agency is deploying neuro-surveillance against citizens, or whether a clinical trial is concealing adverse neurological effects, its pronouncements are hollow.

The challenge is formidable. Covert neurotechnology use does not announce itself. It leaves no shell casings, no fingerprints, no visible wounds. The evidence exists in the electromagnetic spectrum, in biological markers, in patterns that conventional forensic training does not teach. Technical audit mechanisms must therefore be designed from the ground up to detect what is designed to remain undetected.

Spectrum Monitoring and Environmental Forensics

The most direct technical audit mechanism is radio frequency (RF) spectrum monitoring. Every device that emits electromagnetic radiation — and virtually all neurotechnology devices do — leaves a signature in the spectrum. The challenge is distinguishing legitimate emissions from covert ones, and detecting emissions that are intentionally masked or pulsed at irregular intervals.

The source material from the TI One Voice platform proposes a concrete mechanism: a specialized technical assessment unit within the intelligence oversight body, staffed with engineers and physicists with appropriate security clearances, capable of independently auditing electromagnetic emissions from suspected government installations. This unit would need:

Fixed monitoring stations in urban areas, capable of continuous spectrum analysis from DC to millimeter-wave frequencies, with automated alerting for anomalous emissions patterns.

Mobile monitoring teams equipped with portable spectrum analyzers, directional antennas, and signal classification software, deployable in response to citizen complaints within 48 hours.

Forensic recording capability for evidentiary purposes, with chain-of-custody protocols that meet criminal court standards.

Reference libraries of known neurotechnology device signatures, compiled from published research, patent filings, and voluntary manufacturer disclosures.

The Belgian legislative reform proposals in the source material call for a reversal of evidentiary burden regarding the technical feasibility of alleged devices once a victim can demonstrate a pattern of symptoms consistent with published bioeffects literature. This is a

critical innovation: it means that a citizen who presents documented symptoms — tinnitus, sleep disruption, cognitive fog, unexplained burns — can trigger an audit obligation, rather than bearing the impossible burden of proving exactly which device is causing the harm.

Device Registration and Supply Chain Tracking

Verification cannot rely solely on reactive monitoring. It requires proactive knowledge of what devices exist, who possesses them, and where they are deployed.

The source material proposes mandatory registration of all directed-energy devices above defined power thresholds with a new EU-level registry. This is modeled on existing frameworks for dual-use goods, radioactive materials, and controlled chemicals. The registry would include:

Manufacturer declarations of device specifications, intended uses, and safety testing results.

Chain-of-custody records from manufacture through sale, transfer, and disposal, auditable by the oversight body.

Annual compliance certifications from all registered device holders, confirming that devices have not been modified, transferred, or used in unauthorized ways.

Whistleblower reporting mechanisms for employees of manufacturers, research laboratories, and government agencies who become aware of unregistered devices or unauthorized uses.

The source material from the OHCHR submission emphasizes the need for "mandatory transparency requirements for neurotechnology research with dual-use potential." This is not about compromising legitimate research or national security; it is about ensuring that the oversight body knows what exists, so that it can distinguish between authorized and unauthorized activities.

Clinical Trials Oversight

The medical applications of neurotechnology — deep brain stimulation, closed-loop systems, brain-computer interfaces — provide a legitimate pathway for development that can also serve as a cover for dual-use applications. Clinical trials oversight must therefore be designed to detect not only safety violations but also diversion of technology for non-medical purposes.

The source material's proposed Belgian legislation includes a requirement that healthcare professionals consider environmental and technological causes before issuing psychiatric diagnoses for symptom profiles consistent with directed-energy exposure. This has a parallel in clinical trials oversight: when a trial participant reports symptoms consistent with neurotechnology side effects, the protocol must include investigation of whether the device is being used within approved parameters.

Key audit mechanisms for clinical trials include:

Independent data safety monitoring boards with access to raw device telemetry, not just investigator-reported summaries.

Unannounced inspections of trial sites, including physical inspection of devices and comparison of device logs against reported usage.

Mandatory adverse event reporting to a central registry, with automated pattern detection for signals that might indicate misuse.

Post-trial device tracking to ensure that devices are not transferred to unapproved users or applications.

Whistleblower Protections and Secure Reporting Channels

The most powerful audit mechanism is not technical but human. The people who know what is happening — engineers who see a device being modified for unauthorized use, technicians who observe unusual deployment patterns, security personnel who are asked to conceal evidence — must have safe channels to report what they know.

The source material repeatedly emphasizes the need for whistleblower protections for researchers and engineers exposing weaponization. These protections must be:

Comprehensive, covering all forms of neurotechnology work, including classified programs, with exceptions only for information that would directly compromise ongoing operations — and even then, the whistleblower should be able to report to an independent oversight body with appropriate security clearances.

Enforceable, with criminal penalties for retaliation and a private right of action for whistleblowers who suffer adverse employment actions.

Accessible, including anonymous reporting channels, legal representation funded by the oversight body, and protection from defamation lawsuits when reports are made in good faith.

The source material's proposal for a citizen complaint pathway within the Comité R (Belgium's intelligence oversight body) with statutory response deadlines — 90 days for acknowledgement, 12 months for preliminary findings — provides a model for how whistleblower and citizen reports can be processed with urgency and accountability.

Forensic Protocols and Evidentiary Standards

Technical audits are meaningless if their results cannot be used in legal proceedings. The source material from the OHCHR submission calls for "standardized protocols for investigating complaints of technological harassment" and "forensic methodologies to detect and document neurotechnological attacks."

These protocols must address:

Signal detection and classification: How to determine whether a detected RF emission is from a neurotechnology device, and if so, what type and capability.

Biological evidence collection: How to document symptoms in a manner that is medically rigorous and legally admissible, including standardized symptom questionnaires, photographic documentation of physical effects, and biomarker analysis.

Temporal correlation: How to establish that symptoms correlate with device activation, using device logs, environmental monitoring data, and victim diaries.

Source attribution: How to trace a detected emission to a specific device and operator, using signal fingerprinting, triangulation, and chain-of-evidence documentation.

The proposed Belgian legislation's reversal of evidentiary burden is essential here. Without it, victims face an impossible standard: they must prove not only that they are suffering, but that a specific device operated by a specific actor is causing the suffering. With the burden shift, a documented pattern of symptoms consistent with published bioeffects literature creates a presumption that requires the alleged perpetrator to demonstrate that no such device was used.

The Cost of Verification

Establishing these technical audit mechanisms requires investment. Spectrum monitoring equipment, forensic laboratories, trained personnel, and secure reporting systems all cost money. The source material estimates that a fully functional national oversight body with technical audit capacity would cost €5-15 million annually for a country of moderate size. For the European-level registry and coordination mechanism, the costs would be higher — perhaps €50-100 million annually.

But these costs must be weighed against the cost of inaction. The source material documents a human rights crisis in which victims are systematically misdiagnosed, their complaints dismissed, their suffering invisible to legal systems. The cost of verification is the cost of making the invisible visible — and of providing the evidentiary foundation for accountability.

The alternative is a world in which neurotechnology operates without oversight, where covert programs continue undetected, and where the victims of these programs are left without recognition, support, or justice. That cost, measured in human lives and mental integrity, is far higher than any investment in technical audit mechanisms.

CHAPTER 72

Whistleblower Protection in the Neurotechnology Sector

Effective oversight of neurotechnology depends on the willingness of insiders to report wrongdoing. Engineers, scientists, executives, and government officials with knowledge of harmful neurotechnology deployments are essential to accountability. Yet without robust legal protections, those insiders face career destruction and personal risk for speaking. Whistleblower protection is therefore not a peripheral concern but a structural prerequisite for neuro-rights enforcement.

Why Neurotechnology Whistleblowing Is Distinctive

Whistleblowing in the neurotechnology sector presents distinctive challenges that general whistleblower protections do not adequately address.

The technical complexity of the field means that a whistleblower's account often requires expert evaluation to be intelligible to non-technical audiences. A claim that a specific neural decoder has been deployed against civilians, or that a particular directed-energy system has been used in a particular operation, may be true and verifiable, yet incomprehensible without substantial domain expertise. Whistleblowers face the problem of being technically right but rhetorically unconvincing.

The classification regime that surrounds substantial portions of neurotechnology research — particularly defence-related research — creates legal exposure for whistleblowers that does not apply in civilian sectors. Disclosure of classified information can carry criminal penalties including imprisonment, regardless of the public interest served by the disclosure.

The credibility gap created by the Targeted Individual phenomenon affects all whistleblowing in the field. A potential whistleblower considering disclosure must reckon with the likelihood of being associated with a population that institutions have systematically dismissed. The reputational risk of being labelled a "conspiracy theorist" is substantial and may deter even genuine, well-documented disclosure.

The Inadequacy of Existing Frameworks

Existing whistleblower protection frameworks address some but not all of these challenges. The EU Whistleblower Directive (2019/1937), implemented in member state law from 2021, provides protection for disclosure of violations of EU law, including data protection law. The US Whistleblower Protection Act protects federal employees who disclose wrongdoing through prescribed channels. Industry-specific frameworks exist in finance, healthcare, and some defence contexts.

These frameworks share several weaknesses when applied to neurotechnology:

- Limited subject-matter scope. Many frameworks protect disclosure of specific categories of wrongdoing (fraud, regulatory violation, public health) that may not clearly include neurotechnology abuses.
- Channel requirements. Many frameworks require disclosure through specific channels (internal compliance, designated regulators) before external disclosure attracts protection. For neurotechnology, such channels may not exist or may be controlled by the very entities engaged in the wrongdoing.
- National security exclusions. Most frameworks contain national security exclusions that exempt the deployments of greatest concern.
- Burden of proof on the whistleblower. Where retaliation is alleged, the whistleblower typically bears the burden of demonstrating causation, a difficult task in technical industries with multiple grounds for adverse employment action.
- Inadequate remedy. Reinstatement and lost wages, while important, do not compensate for the career destruction and personal cost typically suffered by whistleblowers.

The Required Reform

A whistleblower protection framework adequate to the neurotechnology sector requires several specific provisions:

Explicit subject-matter coverage. The framework should explicitly protect disclosure of neurotechnology-related wrongdoing, including non-consensual deployment, regulatory non-compliance, safety failures, and human rights abuses. The explicit reference closes the interpretive ambiguity that otherwise leaves whistleblowers uncertain whether their disclosure attracts protection.

Multiple authorised channels. Disclosure should be protected through multiple channels, including independent regulators, parliamentary oversight bodies, accredited civil society organisations, and — where other channels fail — journalists. The proliferation of channels prevents capture of any single channel by the entities being reported on.

Modified national security treatment. Where national security classification covers neurotechnology programmes, special procedures should apply. Disclosure to authorised parliamentary oversight bodies should be protected. Disclosure to an independent inspector general with appropriate clearances should be protected. Public disclosure, while remaining subject to classification law, should attract mitigated penalties where the public interest served substantially outweighs the harm of disclosure.

Burden shifting on retaliation. Where adverse employment action follows protected disclosure within a defined period (typically one to three years), the burden of demonstrating non-retaliatory cause should shift to the employer. This procedural reform substantially strengthens enforcement.

Substantive remedy. Remedy for retaliation should include reinstatement (where desired), full back pay and benefits, compensation for career damage including projected future losses, and substantial punitive damages where retaliation is deliberate. Anonymity

protection through the litigation process should be available where the whistleblower's safety so requires.

Affirmative support. Beyond protection from retaliation, whistleblowers require affirmative support: secure communication channels, legal representation funds, mental health support, and — where necessary — protected relocation. A whistleblower protection framework that provides only after-the-fact remedies leaves the whistleblower to bear the full cost of disclosure in real time.

International Coordination

Whistleblower protection in the neurotechnology sector has an essential international dimension. A whistleblower disclosing in one jurisdiction may be subject to extradition, harassment, or surveillance from another. The most prominent national security whistleblowers of the past decade — Manning, Snowden, Assange — have demonstrated the international reach of retaliation.

International protection mechanisms should include:

- Refugee status recognition for whistleblowers fleeing retaliation, with clear standards for assessing the bona fides of their disclosures.
- Mutual legal assistance treaty provisions that exclude whistleblower extradition for disclosures attracting protection in the requested state.
- Diplomatic protection where whistleblowers face state-level retaliation, including consular access, sanctuary provisions in diplomatic missions, and — in extreme cases — diplomatic visas.
- A trusted intermediary infrastructure capable of receiving disclosures across borders, providing initial verification, and routing to appropriate oversight bodies.

The Cultural Dimension

Legal protection, however necessary, is not sufficient. Effective whistleblowing also requires a cultural environment in which disclosure is recognised as legitimate and valuable. The technology sector, including neurotechnology, has historically exhibited a culture in which loyalty to the employer takes precedence over public-interest disclosure. The "move fast and break things" ethos that characterises much of contemporary technology culture is structurally hostile to whistleblowing.

Shifting this culture requires sustained effort: professional codes of conduct that explicitly support whistleblowing, professional society support for members who blow the whistle, recognition of significant whistleblowers through awards and public honours, and inclusion of whistleblower case studies in technical education.

The goal is a professional culture in which an engineer or scientist who encounters wrongdoing recognises a positive duty to disclose, supported by their professional community, protected by their legal system, and respected for their action by the broader public. The achievement of such a culture is among the most consequential factors in determining whether neurotechnology develops in a human-rights-respecting direction.

PART XIV

The Future — 3, 5, 10, and 20 Years Ahead

The trajectory of neurotechnology is not a matter of speculation — it is already inscribed in the market data, patent filings, and regulatory gaps documented in the UN report. The industry has grown from $331 million to $7.3 billion in just ten years. Nearly half of all worldwide neurotechnology patent applications are registered in the United States. Consumer devices are already on the market for well-being, education, and workplace monitoring, despite lacking proven efficacy or long-term safety data.

The future is not a single path but a set of branching possibilities, shaped by the decisions made in the next decade. The UN report is clear: "Where areas not sufficiently covered by existing standards do emerge, action is needed to prevent violations, misuses and abuses." The question is whether that action will come in time.

Three Years (2029): The Tipping Point

The near term is defined by the convergence of three forces: the rapid commercialization of non-invasive neurotechnology, the first major neurodata scandals, and the initial legislative responses.

Consumer-grade EEG headsets and biosensing earbuds are already being marketed for concentration monitoring, sleep tracking, and cognitive performance. The UN report notes that "the sensor systems of the new generation of wireless earbuds include biosensing electrodes, enabling the measurement of various biosignals, including brain activity." By

2029, these devices will be ubiquitous — worn by millions who do not understand that their neural data is being collected, analyzed, and potentially sold.

The first major scandal will involve a company collecting neurodata without meaningful consent and using it for behavioral manipulation. The UN report warns that "data can be sold to third parties without the knowledge or authorization of users" and that "consumer neurotechnology companies are already collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future." When that data becomes decipherable — when companies can infer sexual orientation, personality traits, cognitive performance, intentions, beliefs, and emotions from previously recorded neural signals — the scandal will be global.

Legislative responses will begin in earnest. Following Chile's constitutional reform and the state-level neurodata protections in Colorado and California, other jurisdictions will move to classify neural data as sensitive personal information requiring explicit consent for collection and processing. The UN report's call for "context-specific standards to protect against the non-consensual intrusion of third parties" will begin to be implemented.

Five Years (2031): The Regulatory Race

By 2031, the pressure to regulate will be intense. The UN report identifies five rights particularly at risk: freedom of thought, privacy, personal integrity, health, and the prohibition of torture. Each of these will be tested by specific incidents.

The freedom of thought — described in the report as an absolute right that "does not allow for any interference whatsoever" — will face its first major challenge. As neuromarketing techniques mature and are deployed at scale in political campaigns, the report's warning that "they can be successfully used to condition the forming of opinions, as well as influencing an individual's decision-making processes" will become a demonstrated reality. The first election in which neurotechnology-based behavioral manipulation is proven to have influenced outcomes will trigger a constitutional crisis in the affected country.

The right to personal integrity will be tested by the first documented cases of mental harm from consumer neurotechnology. The UN report notes that "the potential long-term effects of using neurotechnologies are largely unknown" and that "in the absence of a medical need and without being supported by extensive trials and scientific evidence, devices may become harmful for mental health." Lawsuits will follow, and the industry will face its first major liability crisis.

International frameworks will begin to take shape. The UN report's recommendation that "developing specific policies, legislation and instruments" and "reinforcing the institutional structure" are necessary steps will be taken up by coalitions of states. A framework convention on neurotechnology ethics, building on the 2025 UNESCO Recommendation, will enter negotiation.

Ten Years (2036): The Neuronal Divide

The medium-term projections are the most consequential. By 2036, the "neuronal divide" — the gap between those who can access cognitive enhancement and those who cannot — will be a structural feature of society.

The UN report addresses this directly: "Access becomes a key element of the right to health. States should then grant access, without discrimination." But the report also notes that "cutting-edge treatments, such as implants controlled by brain-computer interface software, are often costly and require highly specialized surgeries, which may not always be covered by public health." The result will be a two-tier system: the cognitively enhanced elite and the unenhanced majority, with the gap widening each year.

The first neuro-crimes will be prosecuted. The UN report warns that "some proposed applications of neurotechnologies in the area of justice and national security may lead to the criminalization of mere thought." By 2036, the first case of a person convicted based on neural evidence — or punished for unexpressed intentions detected through neuroimaging — will reach an international human rights body.

The military dimension will also matureThe military dimension will also mature. The UN report flags particular concern about "coercive uses of neurotechnologies to augment resilience on the battlefield (i.e. to reduce emotions, empathy or the need for sleep)" and recommends they "should be prohibited." By 2036, the first documented use of neurotechnology as a weapon in an armed conflict will be reported, triggering emergency sessions of the UN Security Council and accelerating calls for a global ban modeled on the Chemical Weapons Convention.

Twenty Years (2046): The Transhuman Era

The far-term projections are the most difficult to predict with confidence, but the UN report provides the framework. It notes that neurotechnologies "challenge the foundations of the human rights system and can be used in ways that may erode democracy and the rule of law." By 2046, those foundations will be tested to their breaking point.

The convergence of neurotechnology with artificial intelligence will produce systems capable of continuous, bidirectional communication between human brains and digital networks. The UN report describes how "with the support of artificial intelligence and machine learning, allowing for the processing and analysis of large sets of neurodata, the opportunities and uses of neurotechnologies are multiplying." By 2046, these systems will be mature enough to enable what the report calls "human-machine teaming" — but also to enable pervasive neural surveillance.

The concept of mental privacy as traditionally understood will be under existential threat. The UN report's framing of the *forum internum* — "the inner, psychological realm of the person, where thoughts and convictions are formed" — as an absolute right requiring unconditional protection will be tested by technologies that can access that realm without consent or awareness. The report's warning that neurotechnologies "bypass the individual's

conscious control or awareness" and "enable non-consensual external access to thoughts, emotions and mental states" will have become a routine reality for millions.

The most profound question will be about human identity itself. The UN report notes that neurotechnologies can affect "the constitutive elements of the person: the individual's personhood, mental capacities and personality" and asks whether their application "contributes to the objectification and instrumentalization of individuals." By 2046, this will no longer be a theoretical question. The first generation of humans with permanent neural implants will be adults, and the first cohort of children raised with neurotechnology-enhanced education will be entering the workforce. The question of what it means to be human — and whether the *forum internum* can survive the technologies designed to penetrate it — will be the defining ethical challenge of the era.

CHAPTER 73

3 Years From Now (2029) — The Tipping Point

The year 2029 will be remembered as the moment when neurotechnology crossed from niche curiosity into mainstream reality — and when the first major cracks appeared in the legal and ethical frameworks meant to contain it.

Consumer BCIs Go Mainstream

By 2029, non-invasive brain-computer interfaces will have entered the consumer market at scale. Devices resembling sleek headbands, earbuds, and even baseball caps will be sold in electronics stores alongside smartwatches and fitness trackers. These devices, using electroencephalography (EEG) and functional near-infrared spectroscopy (fNIRS), will offer features like attention monitoring for students, stress tracking for workers, and mood optimization for general wellness.

The market that was valued at $7.3 billion in the early 2020s will have more than doubled. Companies like Kernel, NeuroSky, and a new generation of startups will have shipped millions of units. The technology will be marketed as harmless — a tool for self-improvement, productivity, and mental health. But the data these devices collect will be far more intimate than anything previously captured by consumer technology.

The UN Human Rights Council's 2024 report warned precisely about this scenario: "neurotechnologies can be used to decode neural signals and infer mental states, including emotions, intentions, and even thoughts." By 2029, this will no longer be a laboratory capability. It will be happening in millions of homes, schools, and workplaces — often without users fully understanding what they have consented to.

The First Neurodata Scandals

The tipping point will come when the first major neurodata scandal breaks. A consumer BCI company — likely one of the market leaders — will be discovered to have been selling or sharing neural data without meaningful consent. The details will be damning: mood data sold to insurance companies to adjust premiums, attention metrics shared with employers for performance reviews, and emotional state data used by advertisers to target vulnerable consumers at precisely the right moment.

The scandal will trigger a cascade of revelations. Investigative journalists will uncover that several companies have been collecting far more data than their privacy policies disclosed — including data that can be used to infer political beliefs, sexual orientation, and mental health conditions. The UN report's warning that neurotechnology "bypasses the individual's conscious control or awareness" will be proven in the most public way possible.

The public reaction will be swift and angry. Class-action lawsuits will be filed in multiple jurisdictions. Regulators will launch investigations. The term "neurodata scandal" will enter the lexicon, joining the ranks of Cambridge Analytica and other data privacy catastrophes.

Initial Legislation in Major Economies

The scandals will accelerate legislative action that had been stalled by industry lobbying and political inertia. The European Union, already having passed the AI Act in 2024, will move to close the gaps that the Act left regarding neurotechnology. As the UN report noted, "there are no discrete human rights policies to limit the misuse and abuse of neurotechnology" — and by 2029, governments will be scrambling to create them.

Colorado and California, which amended their privacy laws in 2024 to protect neural data as sensitive information, will serve as models for other US states. By 2029, at least a dozen more states will have followed suit, creating a patchwork of protections that will increase pressure for federal legislation. The Federal Trade Commission will issue new guidelines specifically addressing neural data collection and use.

In the European Union, the gaps in the AI Act — its failure to address implants, directed energy, and non-invasive mind-reading devices — will be partially addressed through new regulations specifically targeting neurotechnology. The Council of Europe's Framework Convention on AI will be interpreted to require explicit consent for any collection of neural data.

But legislation will remain uneven. Many countries will still lack any protections at all, creating regulatory havens where companies can test and deploy neurotechnology with minimal oversight. The UN report's call for "a global approach to the governance of neurotechnology" will remain aspirational.

Havana Syndrome Recognition

Perhaps the most significant development of 2029 will be the formal recognition of Havana Syndrome as caused by directed energy weapons. As the Naval Postgraduate School thesis "Neurowar Is Here!" documented, these attacks began in 2016 in Havana, Cuba, and have since affected hundreds of US diplomats, intelligence officials, and military personnel across at least 16 countries.

The thesis noted that Dr. James Giordano, a neuropathologist who investigated the Cuba cases, stated in 2018 that "this is intentional, this is directed, this seems to be a beta test of some type of a viable neuroweapon." By 2029, the weight of evidence will be overwhelming. The National Academies of Sciences 2020 report, which suggested directed energy as the most plausible cause, will be vindicated by declassified intelligence and new forensic capabilities.

Formal recognition will come through multiple channels. The US government will officially classify Havana Syndrome as caused by foreign-directed energy weapons. Victims will receive formal recognition and compensation through expanded programs. The World Health Organization will add directed-energy neuroweapon effects to its classification of diseases.

This recognition will have profound implications. It will confirm that neuroweapons are not speculative future threats but present realities. It will validate the experiences of thousands of victims who have been dismissed, gaslit, and marginalized. And it will force the

international community to confront what the UN report called “the weaponization of neurotechnology” — a challenge that will only grow more urgent in the years ahead.

The Tipping Point

By the end of 2029, the landscape will have fundamentally shifted. Consumer BCIs will be in millions of hands. The first major neurodata scandals will have shattered public trust. Legislation will be moving, though unevenly. And the reality of neuroweapons will be officially acknowledged.

The question will no longer be whether neuro-rights matter. The question will be whether the legal and ethical frameworks can be built quickly enough to protect what the UN report called “the constitutive elements of the person: the individual’s personhood, mental capacities and personality.” The tipping point will have arrived — and the world will be watching to see which way it tips.

CHAPTER 74

5 Years From Now (2031) — The Regulatory Race

By 2031, the landscape of neurotechnology governance will have entered what can only be described as a regulatory race — a frantic, uneven, and often contradictory scramble by nations, international bodies, and commercial entities to establish rules for a technology that is already outpacing every attempt to contain it.

Constitutional Amendments Proliferate

Chile's pioneering 2021 constitutional reform will no longer stand alone. By 2031, at least eight to twelve nations will have amended their constitutions or enacted equivalent fundamental laws to protect neuro-rights. These will not be uniform. Some countries will follow Chile's model, embedding protections for mental privacy, personal identity, and free will directly into their founding documents. Others will take narrower approaches, focusing specifically on neural data as a protected category akin to genetic information or biometric data.

The spread will be geographically uneven. European nations, particularly those with strong data protection traditions, will lead the wave. Several Latin American countries, inspired by Chile's example, will follow. A handful of Asian democracies will join, though with significant variations in scope and enforcement. The United States, constrained by its constitutional structure, will not achieve a federal amendment but will see a patchwork of state-level protections — with Colorado, California, and perhaps a dozen other states having enacted neurodata privacy laws by 2031.

Notably absent from this wave will be the major powers most actively developing neurotechnology for military purposes. China, Russia, and the United States — the three nations most invested in cognitive warfare capabilities — will have no constitutional neuro-rights protections. This absence will be telling, reflecting the tension between national security priorities and individual rights that will define the regulatory race.

The First International Framework

The most significant development by 2031 will be the emergence of the first binding international framework for neurotechnology governance. Building on the foundation laid by the 2022 UN Human Rights Council Resolution 51/3, the 2024 UN report, and the 2025 UNESCO Recommendation, a new treaty or binding protocol will be in advanced negotiation.

This framework will not be comprehensive. It will likely focus on the most universally agreed-upon principles: the prohibition of non-consensual neural data collection, the ban on neurotechnology used for torture or cruel treatment, and the requirement of informed consent for any commercial or medical neurotechnology application. The treaty will draw heavily on existing international law, particularly the Biological Weapons Convention and

the Chemical Weapons Convention, but will explicitly address what Dr. James Giordano identified as the critical gap — that "neuroweapons" fall through the cracks of existing treaties.

The negotiation process will be contentious. As documented in the NATO report on cognitive warfare, "neuroscience is not subject to any rules of international law. A neuroweapon that attacks the brain is neither a biological nor a bioweapon." This legal vacuum will be the central battleground. States with advanced neuroweapons programs will resist binding restrictions. States that have experienced neurotechnology abuse — including those whose diplomats suffered Havana Syndrome — will push for strong prohibitions.

The result will be a compromise: a framework convention that establishes general principles but leaves detailed protocols for future negotiation. It will be a start, but critics will rightly note that it lacks the enforcement mechanisms needed to prevent the worst abuses.

Commercial Neural Implants Go Mainstream

The medical neurotechnology market will have transformed dramatically by 2031. Deep brain stimulation for Parkinson's disease and epilepsy, already established by the 2020s, will become more sophisticated with closed-loop systems that adapt in real-time to neural activity. Brain-computer interfaces for paralysis will move from experimental to approved treatments, with devices from companies like Synchron and others receiving regulatory approval in multiple jurisdictions.

But the most consequential development will be the first wave of non-medical neural implants. By 2031, at least one company will have received approval for a neural implant marketed for cognitive enhancement — not for treating disease, but for augmenting memory, attention, or information processing in healthy individuals. This will be the first crack in the wall separating medical treatment from human enhancement.

The implications will be profound. A new class of "enhanced" individuals will emerge — initially limited to the wealthy and well-connected. This will create what the UN report called "a neuronal divide" between the cognitively enhanced and the unenhanced. The ethical questions raised in the 2025 UNESCO Recommendation about "fair access to cognitive enhancement" will become urgent policy debates.

Neuroweapons Incidents Escalate

The recognition of Havana Syndrome as a neuroweapons attack in 2029 will not end the threat. By 2031, there will be documented incidents of neuroweapons use in at least three distinct contexts.

First, diplomatic targeting will continue and expand. The pattern established in Cuba, China, and other locations will be replicated against diplomats from multiple nations. The attacks will become more sophisticated, harder to detect, and more varied in their effects. Some will cause physical symptoms; others will be purely cognitive — inducing confusion, suggestibility, or memory disruption without detectable physiological markers.

Second, the first confirmed use of neuroweapons against military personnel in a combat zone will occur. The Chinese doctrine of "Military Brain Science," which explicitly includes "injuring the brain" and "interfering with the brain" as operational capabilities, will be demonstrated in a conflict scenario. As the NATO report notes, Chinese strategists have defined the cognitive domain as "the battlefield for conducting ideological penetration" and have identified technologies that "cause brain dysfunction and a loss of control with 'smokeless' methods."

Third, the use of neurotechnology against civilian populations will be documented for the first time in a human rights context. The OHCHR call for input will receive testimonies from individuals in multiple countries reporting systematic technological targeting. As one submission will state, "these applications represent gross violations of: right to mental integrity and privacy, freedom from torture and cruel treatment, right to health and bodily autonomy, freedom of thought and conscience."

The Regulatory Gap Widens

Despite the constitutional amendments, the international framework negotiations, and the growing awareness of neuroweapons, the regulatory gap will actually widen by 2031. The pace of technological development will outstrip the pace of lawmaking. Consumer BCIs will be capable of decoding increasingly complex mental states — not just attention or relaxation, but specific thoughts, emotions, and even internal speech. AI systems will analyze neural data with unprecedented accuracy, enabling applications that the drafters of the 2025 UNESCO Recommendation could not have anticipated.

The most dangerous gap will remain in the military domain. As the NATO report makes clear, "cognitive warfare" is not subject to existing arms control treaties. The convergence of NBIC technologies — nanotechnology, biotechnology, information technology, and cognitive science — will create capabilities that blur the line between weapons, surveillance tools, and information operations. The regulatory race will be, at its core, a race between those who seek to constrain these technologies and those who seek to weaponize them.

By the end of 2031, the outcome of this race will still be uncertain. But the stakes will be clear: the decisions made in the coming years will determine whether neurotechnology becomes a tool for human liberation or the most sophisticated instrument of control ever devised.

CHAPTER 75

10 Years From Now (2036) — The Neuronal Divide

The Two-Tier Brain

By 2036, the predictions made in the 2020s have crystallized into an uncomfortable reality. The "neuronal divide" is no longer a speculative concept — it is a measurable social phenomenon, as real as the digital divide of the early 2000s or the literacy gap of previous centuries. On one side stand the cognitively enhanced: those who can afford neural implants, nootropic regimens, and brain-computer interface subscriptions that boost memory, accelerate learning, and optimize decision-making. On the other side stand the unenhanced majority: those whose brains remain as they were born, increasingly disadvantaged in education, employment, and social mobility.

The market projections from 2024 — $7.3 billion and growing — have proven conservative. By 2036, the neurotechnology industry is a trillion-dollar ecosystem spanning medical devices, consumer wearables, cognitive enhancement services, and neural data markets. Companies like Neuralink and Synchron have delivered on their early promises: commercial neural implants are available in dozens of countries, offering everything from cursor control for paralyzed patients to memory augmentation for healthy professionals. The cost, however, places these technologies firmly in the realm of the wealthy. A basic neural implant with cognitive enhancement capabilities costs approximately $50,000, with annual subscription fees for software updates and cloud-based AI processing adding another $10,000.

The Haves and Have-Nots

The consequences are stark. In elite universities, enhanced students complete degrees in half the time, retaining information with near-perfect accuracy. In competitive industries, enhanced workers process data at speeds unattainable by their unenhanced colleagues. The insurance industry has begun offering premium discounts to individuals with neural implants that monitor stress, fatigue, and cognitive performance — creating a financial incentive for enhancement that further widens the gap.

As the NATO report on cognitive warfare presciently noted, "the massive explosion of behavioural data made available by the advent of social media has empowered researchers to make significant advances in the understanding of the dynamics of individuals and large groups." By 2036, this understanding has been weaponized not only by militaries but by corporations and governments. Neural data — the patterns of thought, emotion, and attention captured by consumer BCIs — has become the most valuable commodity on earth, traded on markets that operate with minimal regulation.

AI-Neural Integration

The convergence of artificial intelligence and neurotechnology has reached a level that seemed impossible a decade earlier. AI systems now serve as the "operating system" for

neural implants, decoding neural signals in real-time, predicting mental states, and optimizing brain-computer interfaces for individual users. The “centaur model” — human-machine teaming — has moved from military research programs to everyday life. Enhanced individuals can query databases with their thoughts, communicate silently through neural links, and access AI-powered cognitive assistants that augment their reasoning.

But this integration comes at a cost. The AI systems that enable these capabilities also have access to the most intimate data ever generated: the firing patterns of individual neurons, the electrical signatures of emotions, the neural correlates of memories and decisions. As one submission to the OHCHR from 2026 warned, “these applications represent gross violations of: right to mental integrity and privacy, freedom from torture and cruel treatment, right to health and bodily autonomy, freedom of thought and conscience.” By 2036, these violations are no longer hypothetical. Multiple cases have been documented of employers demanding access to employees’ neural data, of insurance companies adjusting premiums based on brain activity patterns, and of governments using neural surveillance in criminal investigations.

The First Neuro-Crimes

The legal system has begun to catch up, but slowly. The first neuro-crimes have been prosecuted: cases of unauthorized neural data access, of “brainjacking” — the remote hijacking of neural implants — and of cognitive manipulation through directed energy. The OHCHR’s call for input, which seemed radical in 2026, now appears prescient. The UN report on neurotechnology and human rights (A/HRC/57/61) has become a foundational document, cited in courtrooms and legislative chambers around the world.

Yet the justice gap remains. As one victim’s testimony from the OHCHR process stated, “current transitional justice mechanisms are ill-equipped to address crimes that leave no physical evidence visible to conventional forensics, are perpetrated remotely across jurisdictions, employ technologies officially denied or classified, and create isolation through systematic gaslighting by institutions.” By 2036, this description applies to thousands of victims worldwide.

Global Treaty Negotiations

The most significant development of the decade has been the initiation of formal negotiations for a global treaty on neurotechnology. Modeled on the Biological Weapons Convention and the Chemical Weapons Convention, the proposed treaty seeks to ban neuroweapons, regulate neural data collection, and establish an international oversight mechanism. The negotiations, however, are fraught with difficulty. As the NATO report noted, “neuroscience is not subject to any rules of international law. A neuroweapon that attacks the brain is neither a biological nor a bioweapon. The two existing United Nations treaties — the Biological Weapons Convention (BWC) and the Chemical Weapons Convention (CWC) — do not contain any provisions for neuroweapons.”

The major powers — the United States, China, Russia, and the European Union — have taken divergent positions. China, which has defined “Military Brain Science” as a priority research

area and outlined a comprehensive taxonomy of brain functions including “injuring the brain,” “interfering with the brain,” and “arming the brain,” has resisted restrictions on military applications. The United States, with its extensive DARPA programs and commercial neurotechnology sector, has pushed for a framework that protects industry interests. The European Union, building on its AI Act and the Council of Europe Framework Convention, has advocated for strong human rights protections.

The Unresolved Question

By 2036, the central question posed by this guide remains unanswered: will neurotechnology become a tool for liberation or oppression? The answer depends on the outcome of the treaty negotiations, the evolution of the regulatory frameworks, and the choices made by citizens, legislators, and the global community. The neuronal divide is real, but it is not inevitable. The battle for the last private place continues.

CHAPTER 76

20 Years From Now (2046) — The Transhuman Era

By 2046, the trajectory that was visible in 2026 has reached its logical conclusion. The convergence of neurotechnology, artificial intelligence, and biotechnology has produced a world that the framers of the first neuro-rights laws could scarcely have imagined. The transhuman era is no longer a philosophical speculation — it is a lived reality for millions, and a contested future for billions more.

Human-AI Symbiosis

The most profound transformation of the transhuman era is the widespread integration of human cognition with artificial intelligence. As foreseen by the transhumanist philosophers of the early 21st century, the boundary between human and machine has become permeable. Neural implants that were experimental in the 2020s — Neuralink's N3 device, Synchron's Stentrode, and their successors — have evolved into seamless bidirectional interfaces capable of real-time communication between the human brain and cloud-based AI systems.

The vision articulated in *The Transhumanist Reader* — of "substrate-independent minds" and "whole brain emulation" — has been partially realized. By 2046, individuals with neural implants can access computational resources directly, query databases with thought, and communicate with others through synthetic telepathy. The "centaur model" of human-machine teaming, once a theoretical concept, has become the dominant mode of cognitive work in advanced economies.

Yet this symbiosis has created a new hierarchy. The "neuronal divide" predicted in the 2030s has hardened into a permanent social stratification. Those with access to high-bandwidth neural interfaces — the "enhanced" — enjoy cognitive capabilities that far exceed those of the unenhanced. They process information faster, remember more, and communicate at speeds that make verbal conversation seem glacial. The unenhanced, who cannot afford or choose not to accept neural augmentation, find themselves increasingly marginalized in a society that assumes neural connectivity.

Neural Cloud Connectivity

The neural cloud — a global network of interconnected minds — has emerged as the defining infrastructure of the transhuman era. By 2046, it is possible to share sensory experiences, emotions, and even memories directly between individuals through neural interfaces. The concept of "telempathy" — the direct transmission of emotional states — has moved from science fiction to commercial reality.

The implications for human identity are profound. As the transhumanist philosophers recognized, the question "Who am I?" becomes increasingly difficult to answer when one's thoughts are constantly augmented by AI, when memories can be shared or implanted, and when multiple individuals can occupy the same cognitive space. The "personal identity"

debates that occupied philosophers for centuries have become urgent practical questions for courts, insurers, and families.

The neural cloud has also created unprecedented vulnerabilities. In 2038, a major security breach of a neural cloud provider exposed the private thoughts of over 10 million users. The incident, known as the "Mind Leak," triggered a global crisis of confidence in neural technology. As one victim testified before the UN Human Rights Council: "They didn't steal my credit card number. They stole my memories, my fantasies, my deepest fears. There is no way to change your password for your own mind."

The End of Mental Privacy as We Know It

The most consequential development of the transhuman era is the effective end of mental privacy for those with neural implants. By 2046, the five core neuro-rights — mental privacy, personal identity, free will, fair access to cognitive enhancement, and protection from algorithmic bias — are under assault from multiple directions.

Commercial entities have developed sophisticated neural advertising systems that insert targeted messages directly into the user's stream of consciousness. Employers routinely monitor the neural activity of workers to measure productivity, detect dissent, and predict attrition. Insurance companies require neural data access as a condition of coverage, using predictive algorithms to assess risk based on brain activity patterns.

The legal protections that were fought for in the 2020s and 2030s — Chile's constitutional reform, Colorado and California's neurodata laws, the EU AI Act — have proven insufficient to address the scale of the challenge. As one legal scholar noted in 2045: "We built walls around the fortress of the mind, but the enemy is already inside. The question is no longer whether someone can read your thoughts, but whether you can have any thoughts that are truly your own."

Post-Human Ethics

The transhuman era has forced a fundamental reexamination of ethical frameworks. The concept of "human dignity" — the foundation of the Universal Declaration of Human Rights — has become contested territory. If a person has replaced 80% of their brain with neural implants, are they still human? If a mind has been uploaded to a cloud server, does it have the same rights as a biological person? If a child is born with genetic modifications and neural enhancements from birth, what does "informed consent" mean?

The transhumanist vision of "morphological freedom" — the right to modify one's body and mind without interference — has clashed with the neuro-rights framework's emphasis on protection from coercion. The debate has become increasingly polarized between "enhancement advocates" who argue for maximal individual freedom and "protectionists" who warn of irreversible harm to human identity and social cohesion.

As philosopher Max More wrote in the early 21st century, transhumanism seeks "the continuation and acceleration of the evolution of intelligent life beyond its currently human

form." By 2046, that evolution is underway, but its direction is contested. The question is no longer whether humans will be transformed, but who will control the transformation.

The Need for a New Social Contract

The legal and political frameworks of the 20th and early 21st centuries were designed for a world in which the boundaries of the self were clear and inviolable. The transhuman era has erased those boundaries. The "forum internum" — the inner realm of thought that was once considered sacred and inaccessible — has become a domain of commercial exploitation, state surveillance, and technological manipulation.

The response to this crisis has been fragmented. Some nations have embraced the transhuman future, investing heavily in neural enhancement and treating cognitive augmentation as a national priority. Others have imposed strict regulations, banning non-medical neural implants and criminalizing unauthorized neural data collection. A growing movement of "neo-Luddites" has rejected neural technology entirely, forming communities that maintain the old boundaries between mind and machine.

The most ambitious attempt to establish a new social contract for the transhuman era is the proposed "Global Neuro-Rights Convention," which has been under negotiation since 2041. The convention seeks to establish binding international standards for neural data protection, cognitive enhancement access, and the prohibition of neuroweapons. But the negotiations have been hampered by the same divisions that plagued the earlier treaty discussions: the major powers have conflicting interests, and the technology continues to evolve faster than the law.

As the UN High Commissioner for Human Rights stated in 2045: "We are trying to build a house while the ground beneath us is shifting. The transhuman era has arrived, and we are still arguing about the foundations."

The Unfinished Battle

By 2046, the battle for the last private place is not over — it has entered a new phase. The technologies that were once speculative have become real. The threats that were once hypothetical have become immediate. The choices that were once theoretical have become existential.

The transhuman era offers extraordinary possibilities: the end of disease, the extension of life, the expansion of consciousness, the enhancement of human potential. But it also poses unprecedented risks: the loss of mental privacy, the erosion of personal identity, the creation of new forms of inequality, the weaponization of the mind itself.

The outcome of this battle will determine not only the future of neuro-rights, but the future of what it means to be human. The decisions made in the coming years — by citizens, legislators, and the global community — will shape the transhuman era for generations to come. The last private place is still worth fighting for.

CHAPTER 76-bis

Geopolitical Impact — How Neurotechnology Reshapes World Politics

Part XIV's chapters on the three-, five-, ten- and twenty-year horizons focused on the technical and legal evolution of neurotechnology. This chapter complements that material by tracing what neurotechnology will do to the relations among states, the balance of global power, and the structures of international order. The argument is that neurotechnology will reshape geopolitics across four time horizons in distinct ways, each more consequential than the last, and that the choices made in the near horizon will determine which trajectory the longer horizons follow.

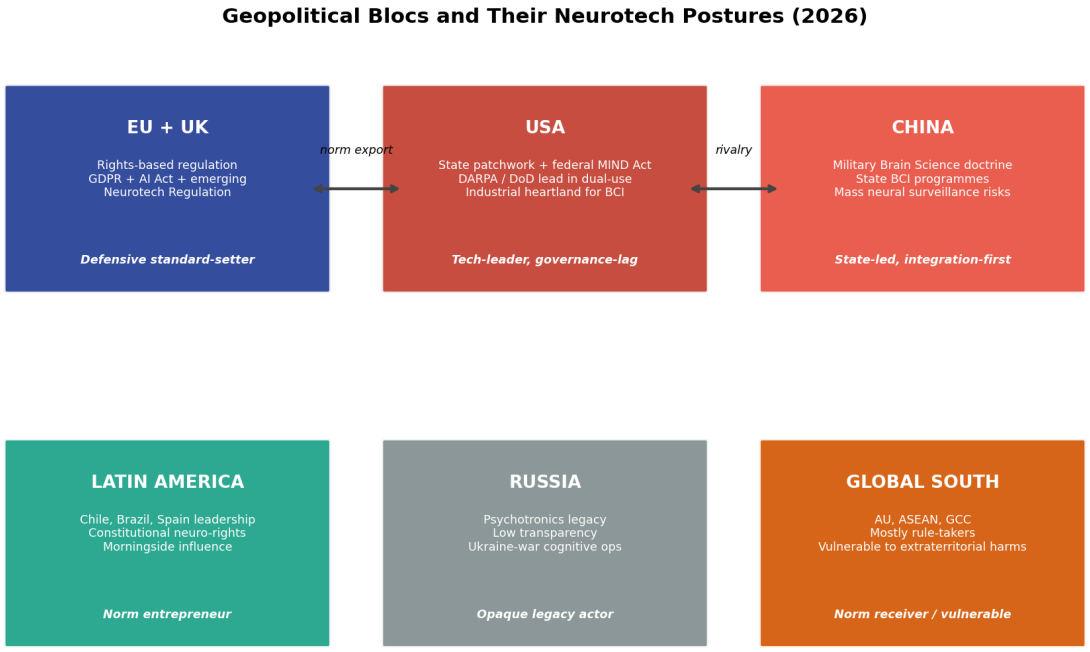


Figure 76bis.1. Geopolitical blocs and their neurotech postures in 2026. The EU is the defensive standard-setter; the US is the tech-leader with a governance lag; China is the state-led, integration-first actor. Latin America is the principal norm entrepreneur; Russia is the opaque legacy actor; the Global South is structurally vulnerable as a rule-receiver.

The Current Configuration

In 2026, the global neurotechnology landscape is organised around six recognisable blocs, each with a distinct posture.

The European Union and the United Kingdom together constitute the world's principal defensive standard-setter. The combination of the GDPR, the AI Act, the Council of Europe AI Convention, and a still-emerging dedicated Neurotechnology Regulation gives the EU bloc the most articulated rights-based framework for neural data, even though it has no constitutional neuro-rights clause. The bloc's strategy is the one it has used successfully for data protection: set the rules, export them through the Brussels Effect, and accept the slower domestic industry growth as the price of normative leadership.

The United States is the global tech-leader with a substantial governance lag. American companies — Neuralink, Synchron, Blackrock, Paradromics, Precision Neuroscience, Kernel — occupy the leading positions across most invasiveness tiers. American defence research

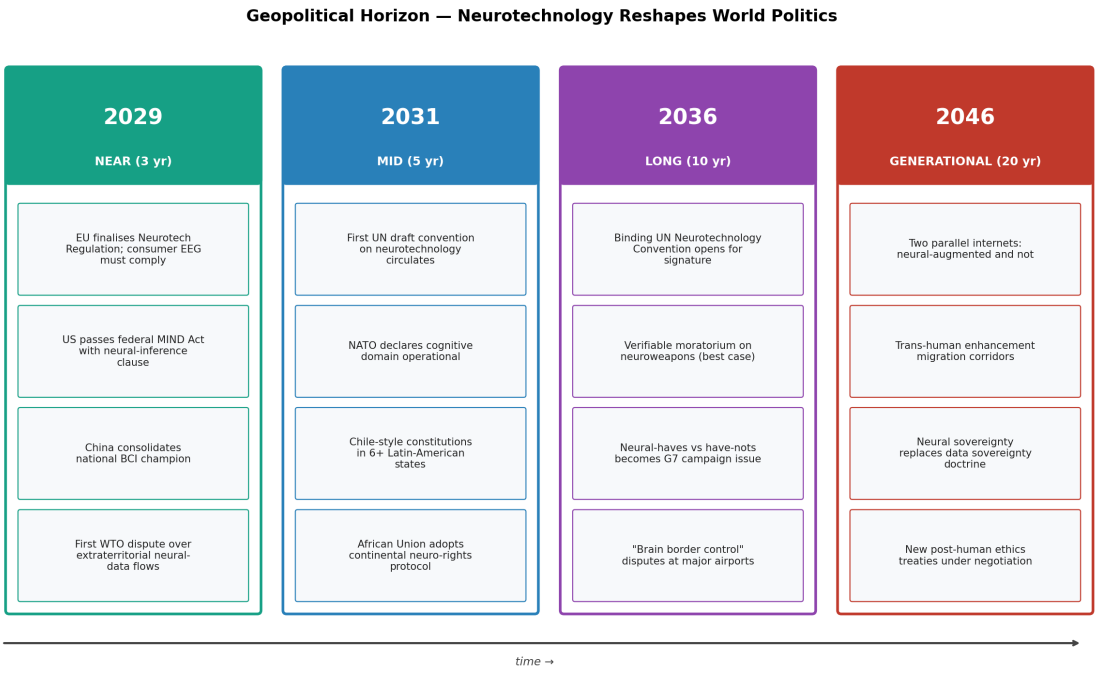
(DARPA, BRAIN Initiative) leads in dual-use capability. But American federal regulation lags badly: no general privacy law, state-by-state patchwork, slow-moving FDA processes for breakthrough categories. The MIND Act of 2025 is the first federal neurodata-specific proposal; it remains pending.

China is the state-led, integration-first actor. The Military Brain Science doctrine (Chapter 29) integrates neurotechnology with surveillance, military operations, and population management in a way no Western state would attempt. Chinese consumer-neurotech firms (BrainCo, others) have aggressive school and workplace deployments. Cross-border data flows from Chinese-manufactured wearables raise extraterritorial concerns for users worldwide. The Chinese posture is the principal model for what authoritarian neurotech looks like at state scale.

Latin America is the principal norm entrepreneur. Chile, Brazil, Spain (a European norm partner), Mexico, Colombia, Argentina and Uruguay together represent the world's most rights-progressive cluster on this topic. The 2021 Chilean constitutional amendment, the 2023 Constitutional Court ruling, the Brazilian amendment, and the Inter-American Declaration of Principles together constitute the doctrinal leadership of the global neuro-rights movement. The bloc has limited industrial capacity but disproportionate norm-setting influence.

Russia is the opaque legacy actor. The Soviet psychotronics legacy (Chapter 28), the constraints of the post-Soviet research ecosystem, the Ukraine war's cognitive-operations dimension, and the general opacity of Russian state research together produce an actor about which the open-source record is limited but the inherited capability is real. Russia's principal influence in the next decade is as a spoiler in international negotiations rather than as a primary actor.

The Global South — comprising the African Union, ASEAN, the Gulf Cooperation Council, South Asian states, and the Pacific Small Island Developing States — is structurally a rule-receiver rather than a rule-maker. The Global South bears most of the externality of neurotechnology developed elsewhere: consumer products designed and regulated for Western markets sold without equivalent protection in non-Western markets; cross-border data flows that concentrate value in foreign servers; extraterritorial application of foreign-state surveillance through the same channels. The Global South's principal interest is in collective bargaining for substantive protection — the case made in Chapter 14.



The Four Horizons

Figure 76bis.2. The geopolitical horizon. Near-term (3-year) developments are highly probable and shape the medium term; long-term (20-year) developments are conditional on choices made now.

Near Horizon — 2029

The near horizon is the period of regulatory consolidation. The European Union finalises a dedicated Neurotechnology Regulation that extends the AI Act's neural provisions and integrates with the GDPR. Consumer EEG products sold in the EU must comply with new labelling, consent and inference-disclosure rules; non-compliant products are blocked. The United States passes the MIND Act with a substantive neural-inference clause; the FTC opens a series of high-profile investigations into consumer-neurotech data practices. China consolidates its national BCI champion through state-directed industrial policy; cross-border data flows from Chinese consumer-neurotech firms become a major bilateral issue with the US and EU. The WTO sees the first dispute over extraterritorial neural-data flows, anticipating the trade-law battles to come.

The principal geopolitical effect of this period is the consolidation of three distinct regulatory blocs. The EU, the US (federal plus pioneer states), and China each have substantially different rules. The blocs' rules diverge enough to constitute a fragmented international regime; they overlap enough to allow some interoperability; companies operating across blocs face significant compliance costs. The pattern resembles the early-2020s data-protection landscape, with the same kinds of friction.

Medium Horizon — 2031

The medium horizon is the period of international-framework construction. The first United Nations draft convention on neurotechnology circulates among member states; negotiations are slow but the document is now part of the international diplomatic agenda. NATO formally declares the cognitive domain operational, with rules of engagement, defence postures, and a coordinating cell. Six or more Latin American states have adopted Chile-style constitutional neuro-rights clauses, anchoring the regional doctrine. The African Union adopts a continental neuro-rights protocol modelled on the Malabo Convention but updated for inference protection.

Geopolitically, this period is when neurotechnology becomes a recognised dimension of international relations. The G7, G20, and BRICS forums all add neurotechnology to their agendas. Bilateral neurotechnology cooperation agreements emerge between major blocs. The principal tension is between the Western (rights-based) and the Chinese (state-led) models, with the Global South positioned as the swing constituency in international forums.

Long Horizon — 2036

The long horizon is the period of treaty consolidation or fragmentation, depending on choices made in the near and medium horizons. In the optimistic scenario, the binding UN Neurotechnology Convention opens for signature, with a verification regime modelled on the Chemical Weapons Convention. A verifiable moratorium on neuroweapons becomes politically achievable, although difficult states (Russia, China, possibly the US under specific administrations) resist accession. The neural-haves-versus-neural-have-nots debate becomes a major campaign issue in G7 elections, with implications for industrial and tax policy. "Brain border control" disputes become a recurring issue at major airports as states begin to claim authority to scan or refuse entry to travellers wearing or implanted with neurotechnology.

In the pessimistic scenario, the UN Convention fails or is reduced to a non-binding declaration. Neuroweapons proliferate among middle-power states, and the absence of verification makes a Cold-War-style standoff plausible. The cognitive domain becomes a primary site of state-on-state competition with no agreed rules of engagement. The neural divide between augmented and unaugmented populations within G7 states becomes a domestic political crisis with significant external implications.

Generational Horizon — 2046

The generational horizon is where forecasting becomes increasingly conditional. The most likely structural development is the emergence of two parallel digital ecosystems: one for the neurally augmented (with embedded BCI capability standardised across consumer devices, with neural-direct authentication and payment, with workflows that presuppose BCI input) and one for the unaugmented. The economic, social and political implications of such a split are enormous and largely uncharted.

Geopolitically, several novel patterns become plausible. Trans-human enhancement migration corridors — pathways through which people in jurisdictions that prohibit certain enhancements travel to jurisdictions that permit them — become a recognisable pattern, analogous to current medical-tourism patterns but with deeper implications for citizenship and identity. The concept of "neural sovereignty" begins to replace "data sovereignty" as the principal frame for thinking about state authority over information flows, because the neurally augmented person carries a continuous bi-directional channel between their body and the network. New post-human ethics treaties are under negotiation, attempting to govern the augmented populations that the previous generation's treaties did not anticipate.

In the optimistic scenario, the choices made in the near and medium horizons produce a stable international order in which neurotechnology is governed under binding rules with effective verification, augmentation is broadly available with effective equality-of-access mechanisms, and the post-human ethics treaties are themselves a sign of mature international institutions adapting to new conditions. In the pessimistic scenario, the same period produces a fragmented and increasingly stratified global order in which neurally augmented populations within powerful states are effectively a separate class with separate political weight.

The Bloc Dynamics

EU vs US — Norms and Industry

The EU-US relationship will be the principal axis of Western neurotechnology governance through the medium horizon. The pattern will broadly resemble the data-protection era: the EU sets normative standards through dedicated regulation; American companies comply because the EU market is too large to ignore; the EU framework becomes the de facto global baseline for the rule-of-law bloc; the US debates internally about whether to harmonise upward (matching the EU) or sideways (defending American practices). The principal uncertainty is whether the US passes a comprehensive federal framework or remains a state-by-state patchwork; the latter is the more likely outcome through 2030.

US-China — Rivalry and Decoupling

The US-China dynamic on neurotechnology will be intense, complicated, and increasingly important. The technological rivalry is real: each side wants the leading position in BCI capability, in foundation models for the brain, in supply chains for high-density implants. The governance rivalry is equally real: the Chinese model of state-led integration of neurotechnology with surveillance and military uses is fundamentally incompatible with the Western rights-based model. Decoupling pressures will grow: restrictions on cross-border investment, on export of advanced manufacturing equipment, on data flows from Chinese consumer-neurotech devices. The cognitive-warfare dimension makes the rivalry even more sensitive, because plausibly-deniable operations can produce significant strategic effects without crossing kinetic thresholds.

EU-China — Engagement and Hedging

The EU's relationship with China on neurotechnology will be more nuanced than the US-China rivalry. The EU's economic dependencies make complete decoupling unattractive; the EU's normative commitments make uncritical engagement untenable. The likely pattern is hedging: substantive engagement on standards (where the EU benefits from Chinese cooperation in international forums even when it disagrees with Chinese practice), targeted restriction of high-risk products (Chinese consumer-neurotech with inadequate data-

protection guarantees), and active norm-promotion in third-party jurisdictions (the Global South) where the two blocs compete for influence.

The Latin American Vanguard

Latin America's role through the medium horizon is the principal source of pro-neuro-rights normative leadership in international forums. Chile, Brazil, Spain (as a European norm partner), Mexico, Colombia, Argentina and Uruguay together have a louder voice on this topic than their combined economic weight suggests. Their cohesion through the next decade is the principal variable that determines whether the optimistic 2036 scenario is achievable. If the Latin American cluster fractures (as has happened on other rights issues under various administrations), the pro-rights coalition in international negotiations loses its principal articulator.

The Global South and Multilateral Bargaining

The Global South's leverage in neurotechnology governance is essentially multilateral: at the UN, at the World Health Organization, at the WTO, and at the regional bodies (AU, ASEAN, OAS, GCC). The leverage is real but easily fragmented. The principal Global South interest is in collective bargaining for substantive protection rather than the formal-equality arrangements that have characterised some previous technology-governance regimes. The success of the African Union neuro-rights protocol, anticipated in the medium horizon, will be a major test.

Three Geopolitical Scenarios for 2046

We close this chapter with three explicitly labelled scenarios for the generational horizon. They are scenario-planning, not prediction; the value lies in their internal consistency and in the contrast among them.

Scenario A — The Layered Order (optimistic)

By 2046, the binding UN Neurotechnology Convention has been operational for a decade. Its verification regime, modelled on the OPCW, conducts inspections of facilities suspected of neuroweapons development; member states have ceded sovereignty over specific verification protocols in exchange for the security gains. The Convention has been ratified by all G7 and G20 states, with the principal hold-outs being a small number of authoritarian states that face increasing diplomatic isolation. National neuro-rights frameworks are operational in approximately seventy states; the African Union protocol is implemented in fifty-four members. Augmentation is broadly available in high-income states through public-insurance pathways; the neural divide is significant but managed through redistributive policy. Cognitive warfare is constrained by mutual deterrence and verification. The Latin American vanguard is recognised as the founding principal of the modern regime.

Scenario B — The Fragmented Order (baseline)

By 2046, the UN Convention process has produced a non-binding declaration with limited verification. Three principal regulatory blocs operate: the rights-based bloc led by the EU and Latin America, the state-led bloc led by China, and a flexible-rules American bloc that varies with administration. Cognitive warfare is endemic between the blocs at low intensity; mutual restraint depends on case-by-case calculation rather than treaty constraint. Cross-border neurotechnology trade is heavily regulated, with substantial frictional costs. The neural divide is significant and politicised within states. Augmentation migration corridors exist. International forums produce incremental adjustments without breakthrough.

Scenario C — The Stratified Order (pessimistic)

By 2046, no binding international instrument exists. Neuroweapons capability has proliferated to approximately a dozen states. The cognitive domain is a primary site of state-on-state competition with no rules of engagement; deniable operations against civilian

populations of adversary states are routine. Within high-income states, the neural divide produces a stratified society with separate political weight for the augmented; democratic institutions are under sustained pressure. Cross-border data flows are governed by data-localisation rules in each major bloc; the consumer internet has fragmented. The Latin American vanguard's normative leadership has been overwhelmed by the practical realities of an unrestrained capability race. The window described in this book has closed.
Which scenario actually unfolds will be determined by the policy choices made in the near and medium horizons. The choices are not yet made. The argument of this book is that the layered-order scenario is technically achievable and morally required, and that the policy effort to achieve it must begin now.

PART XV

What Top-Secret Neurotechnology Might Exist?

Informed speculation about classified neurotechnology must begin with a recognition of what is already publicly documented. The historical record shows a continuous trajectory from laboratory research to weaponization that spans more than five decades.

The Historical Foundation

The 1974 HEW report *Individual Rights and the Federal Role in Behavior Modification* documented federal grants for “direct electrical stimulation of the brain” and raised early concerns about the implications of such research. CIA brain implant research has been traced across five decades, as documented in the *Hi-Tech War Crime Report*. These programs established the foundational knowledge that would later be refined into operational systems.

What is known publicly provides a framework for understanding what likely exists in classified programs. The scientific principles are established. The technical challenges have

been addressed, at least partially. The question is not whether such technologies *could* exist, but rather which specific capabilities have been achieved and deployed.

Remote Neural Monitoring

The most frequently reported capability in victim testimony is remote neural monitoring — the ability to read brain activity at a distance without physical contact. The scientific basis for this is well-established. Functional MRI and EEG can decode certain mental states, but these require proximity or physical contact. The classified question is whether these techniques have been extended to remote, covert monitoring.

Binhi's analysis notes that "EM radiation can be easily organized so as to unexpectedly or even secretly expose large groups of people using relatively distant EM sources." The penetrating power of electromagnetic fields makes them ideal for covert applications. The technical literature on "ion cyclotron resonance" effects — where specific magnetic field combinations affect rat behavior, memory, and aggression — suggests that frequency-selective neural manipulation is scientifically feasible.

The *Electronic Mind Control Book* documents that "rats temporarily lose their short-term memory after just one hour of exposure to ICR fields" and that "rats exposed to these magnetic field combinations can be made to learn either more quickly than controls or less, depending on the field combination that is used." These are not speculative effects — they are replicated experimental results.

Synthetic Telepathy

The microwave hearing effect — where pulsed microwaves create the perception of sound inside the head — is a well-documented phenomenon. Binhi notes that "in principle one can modulate a high frequency signal with low frequency patterns that correspond to speech, and even electronically create certain statements."

The Motorola "electronic telepathy" patent demonstrates that corporate research has explored this capability. The question is whether military programs have refined this into operational systems capable of delivering intelligible speech to targeted individuals at distance.

Victim testimony consistently describes "voice-to-skull" harassment — the experience of hearing voices that appear to originate from external sources but are not audible to others. While this phenomenon is often dismissed as psychosis, the technical capability to produce such effects has been demonstrated in controlled experiments.

Memory Manipulation and Behavior Modification

The research documented in the *Electronic Mind Control Book* shows that electromagnetic fields can affect memory, learning, and aggression in animal models. The extension to humans is scientifically plausible, though the ethical implications are profound.

DARPA's official neural interfacing programs — particularly the Next-Generation Non-Surgical Neurotechnology (N3) program — demonstrate military interest in creating neural interfaces for able-bodied soldiers. The stated goal is communication and control, but the same technology could be adapted for manipulation.

The *Hi-Tech War Crime Report* documents a "continuous trajectory from research to weaponization." This trajectory suggests that classified programs have likely achieved capabilities that go well beyond what is publicly acknowledged.

The Verification Challenge

The most troubling aspect of classified neurotechnology is the difficulty of verification. As the UN response document notes, victims "lack formal complaint procedures that recognize technological attacks" and "investigative bodies lack technical expertise to verify neurotechnological harassment."

The technologies that likely exist in classified programs — remote neural monitoring, synthetic telepathy, memory manipulation, behavior modification via directed energy — are kept secret not because they are impossible, but because they represent the ultimate form of covert control. As Binhi observes, "if everyone had access to such technology, it would quickly destabilize a society."

The question is not whether these technologies exist, but whether the international community has the courage to acknowledge them and create the legal frameworks necessary to protect citizens from their abuse.

CHAPTER 77

From MKULTRA to the Present — A History of Secret Brain Research

The history of secret brain research is not a story of fringe science or conspiracy theory. It is a documented history of government-funded programs, spanning decades, that sought to understand and ultimately control the human mind. Understanding this history is essential for grasping the current trajectory of neurotechnology and the urgency of neuro-rights.

The MKULTRA Era: Systematic Mind Control Research

The most well-known program, MKULTRA, was initiated by the Central Intelligence Agency in 1953. Its stated purpose was to develop mind-control techniques for use in interrogation, espionage, and covert operations. The program explored a wide range of methods: drugs (LSD, barbiturates), hypnosis, sensory deprivation, electroshock, and psychological torture.

MKULTRA was not a single project but a network of subprojects involving universities, hospitals, prisons, and private research institutions. Subjects were often unwitting — prisoners, mental patients, and even CIA employees who were given drugs without their knowledge. The program produced thousands of pages of documentation, much of which was destroyed in 1973 on the orders of then-CIA Director Richard Helms.

What MKULTRA revealed was that governments were willing to invest heavily in understanding how to alter human consciousness, memory, and behavior. The program established a template: classify the research, use unwitting subjects, and deny any wrongdoing when exposed.

Project ARTICHOKE and the Search for "Special Interrogation"

Running parallel to MKULTRA was Project ARTICHOKE, also overseen by the CIA and the Office of Naval Research. ARTICHOKE focused on “special interrogation” techniques — methods to extract information from resistant subjects and to induce behavioral changes. The project explored the use of drugs, hypnosis, and electroshock to create amnesia, suggestibility, and compliance.

Declassified documents from ARTICHOKE reveal a fascination with the possibility of “remote control” of human behavior. Researchers sought to identify whether electromagnetic fields or other physical agents could be used to influence brain function at a distance. The project’s findings were classified and remain largely unavailable.

Soviet Psychotronics Research

The Soviet Union conducted extensive research into what it called “psychotronic weapons” — devices that could influence human consciousness through electromagnetic fields. As noted in the *Electronic Mind Control Book*, Soviet researchers explored the use of microwave

radiation to affect brain function, building on the microwave hearing effect discovered in the 1950s.

The most famous incident was the microwave irradiation of the US Embassy in Moscow during the Cold War. For years, the Soviet Union directed microwave beams at the embassy building. The purpose remains unclear — it may have been an attempt to disrupt electronic surveillance, or it may have been an experimental test of biological effects on embassy personnel. The Project Pandora Report, commissioned by the US State Department, came to no clear conclusion.

Soviet research into electromagnetic effects on the brain was extensive and scientifically rigorous. As Binhi documents, Soviet scientists made significant contributions to understanding nonthermal biological effects of electromagnetic fields, including effects on memory, learning, and behavior in animal models.

The Post-Cold War Continuity

After the Cold War, research into brain control did not end — it went deeper underground. The *Electronic Mind Control Book* notes that "military research may contain secret knowledge, or know-how" and that "it is the technology or know-how rather than the scientific concept that usually contains sensitive information that an institution will classify."

DARPA's N3 program (Next-Generation Non-Surgical Neurotechnology) represents the public face of this research. But as the source material makes clear, "there are many reasons to expect that publicly available scientific literature would contain explicit evidence that EM fields and radiations are capable of controlling the human mind, in principle." The question is whether direct studies on EM mind control have been classified.

The Pattern of Denial and Disinformation

A recurring pattern in the history of secret brain research is the use of denial and disinformation to discredit those who report being affected. As noted in the Foreword to Binhi's book, Dr. Robert O. Becker "always maintained that both the US and Russian governments were very much involved in such research. But he also felt that there was a deliberate attempt by authorities to cover up these clandestine activities by means of deliberate disinformation tactics. These included weird and bizarre stories that were leaked to the public for one purpose: to muddy the waters, to make the whole question of mind control seem absolutely unbelievable."

This pattern has been observed in relation to Havana Syndrome, where victims have been dismissed as having mass psychogenic illness or stress-related disorders. It has been observed in relation to targeted individuals who report being subjected to directed-energy attacks. The *Electronic Mind Control Book* notes that "sometimes, it is actually hard to distinguish the truth from conspiracy theory" — a deliberate ambiguity that serves to protect classified programs.

The Present: Where Are We Now?

The history of secret brain research is not history. It is ongoing. The technologies that were speculative in the 1950s — remote neural monitoring, behavior modification via electromagnetic fields, memory manipulation — are now scientifically plausible. As the *Electronic Mind Control Book* concludes, "based on the review of literature made, we deduce which kind of mind reading is scientifically grounded and in what sense we may speak about mind control."

The question is not whether these technologies exist. The question is what legal frameworks are in place to protect citizens from their abuse. And the answer, as this guide has documented, is that the legal frameworks are almost entirely absent.

The battle for the last private place — the human mind — has been fought in secret for decades. It is time to bring it into the light.

CHAPTER 78

Speculative Technologies — What Likely Exists but Is Classified

The history of secret brain research, from MKULTRA to DARPA's N3 program, establishes a clear pattern: technologies that are scientifically plausible are developed in classified settings, often decades before they are acknowledged publicly. The question is not whether remote neural monitoring, synthetic telepathy, memory manipulation, and behavior modification via directed energy exist. The question is which of these capabilities have been operationalized, and at what technical level.

Remote Neural Monitoring at Distance

The scientific foundation for remote neural monitoring rests on the ability to detect and decode the electromagnetic signatures of neural activity. As documented in the *Electronic Mind Control Book*, "the progress in cognitive studies has generated huge scientific, military, and governmental interest in the development and future applications of neurotechnologies." The book notes that "modern functional neuroimaging makes it possible to visualize and measure different brain functions, enabling one to detect the association of activity in a particular portion of the brain with the specific features of cognitive processes and behavior."

The critical question is whether this detection can be performed at a distance, without the subject's knowledge or consent. The *Electronic Mind Control Book* addresses this directly: "EM radiation can be easily organized so as to unexpectedly or even secretly expose large groups of people using relatively distant EM sources. Or, under certain conditions, EM radiation can be focused on a local target."

The technical requirements for remote neural monitoring include:

53. **Sensitive receivers** capable of detecting the extremely weak electromagnetic signals generated by neural activity, which are on the order of picoteslas to femtoteslas — many orders of magnitude weaker than the Earth's magnetic field.
54. **Signal processing algorithms** that can extract meaningful neural patterns from background noise, including the subject's own physiological signals and environmental electromagnetic interference.
55. **Decoding models** that can translate neural patterns into specific thoughts, emotions, or mental states — a challenge the *Electronic Mind Control Book* describes as requiring "a comprehensive computer program" that would "have to replace the brain in its most enigmatic and subtle function, consciousness."

The *Electronic Mind Control Book* is cautious about the feasibility of full mind reading: "It seems impossible even within the physicalist framework." However, this assessment was made in the context of reading *all* thoughts — a far more demanding task than detecting specific mental states, emotional arousal, or the presence of particular cognitive processes. The distinction is critical.

What likely exists in classified settings is not the ability to read every thought, but the ability to:

Detect when a subject is thinking about a specific topic (e.g., a person, place, or event)

Monitor emotional arousal levels and stress responses

Identify the presence of deception or concealed information

Track attention and cognitive load

These capabilities are consistent with the trajectory of research documented in the *Electronic Mind Control Book*, including the calcium efflux experiments by Adey and Bawin, the ion cyclotron resonance studies by Liboff and Thomas, and the behavioral modification experiments on rats showing that "rats temporarily lose their short-term memory after just one hour of exposure to ICR fields."

Synthetic Telepathy (Voice-to-Skull)

Synthetic telepathy — the transmission of speech or other auditory information directly into a person's brain via electromagnetic or ultrasonic means — is perhaps the most well-documented speculative technology. The mechanism is the microwave hearing effect, also known as the Frey effect, in which pulsed microwave radiation is perceived as clicks, buzzing, or even intelligible speech.

The *Electronic Mind Control Book* describes this effect: "It depends on the use of the acoustic response that sometimes occurs when the head is exposed to microwave signals, otherwise called microwave hearing. This microwave-induced acoustic response is a very real thing."

The book's foreword, written by Dr. A.R. Liboff, elaborates: "In principle one can modulate a high frequency signal with low frequency patterns that correspond to speech, and even electronically create certain statements."

What likely exists in classified settings is a refined version of this technology that can:

Transmit intelligible speech directly into the brain without external speakers

Target specific individuals, even in crowds, using phased-array antenna systems

Operate at power levels below those that would cause thermal damage or detectable heating

Maintain plausible deniability by mimicking the symptoms of mental illness

The *Electronic Mind Control Book* notes that "if we were to assume that this was technically feasible, I doubt very much that mind 'take-over' would ever occur. It is more likely that subjects would simply become greatly distracted, driven out of their minds, or even become suicidal."

This assessment is consistent with the experiences reported by targeted individuals, who describe persistent auditory harassment, sleep deprivation, and psychological distress — not the "zombie-like effects" of popular imagination, but a more insidious form of control through disorientation and exhaustion.

Memory Manipulation

The ability to selectively enhance, suppress, or alter memories has been demonstrated in animal studies and is scientifically plausible in humans. The *Electronic Mind Control Book* cites experiments showing that "in all cases tested, rats temporarily lose their short-term memory after just one hour of exposure to ICR fields."

Further experiments showed that "rats exposed to these magnetic field combinations can be made to learn either more quickly than controls or less, depending on the field combination that is used." The book connects these effects to "a key enzyme in the brain, the NMDA glutamate receptor," which "has long been implicated by neurophysiologists in memory and learning."

What likely exists in classified settings is the ability to:

Temporarily disrupt short-term memory formation, creating gaps in a subject's recollection

Enhance or suppress specific memories by targeting the neural circuits involved in memory consolidation

Implant false memories by creating artificial neural patterns that the brain subsequently integrates into its existing memory structures

The *Electronic Mind Control Book* is explicit about the behavioral implications: "These magnetic fields are also capable of producing greater aggressiveness or greater passivity, again as compared to controls. In effect, the way has been pointed towards selective frequency control of behavior."

Behavior Modification via Directed Energy

The most concerning speculative technology is the ability to modify behavior — to make a person more aggressive, more passive, more fearful, or more compliant — through the application of specific electromagnetic frequencies.

The *Electronic Mind Control Book* documents the scientific basis for this capability: "There is good reason to think that these effects on important behavioral functions can also be made to occur in humans. Thus we think it highly feasible that we can induce loss of memory in people, make individuals more fearful or perhaps more willing to take chances, or simply compromise one's ability to make simple decisions."

The book's foreword describes the implications: "There can be no question that electromagnetic effects like these represent extraordinary opportunities for the military, to say nothing of corporate advertising or even political advancement."

What likely exists in classified settings is a suite of directed-energy systems that can:

Induce specific emotional states (fear, calm, aggression, compliance)

Disrupt cognitive functions (attention, decision-making, memory)

Alter perception (visual or auditory hallucinations)

Create physical sensations (pain, pressure, temperature changes)

The *Electronic Mind Control Book* notes that "the underlying research is likely focused on two very different frequency ranges, first, at extremely low frequencies, from just one or two to several hundred hertz, and second, in the gigahertz region and beyond. It is also highly probable that signals encompassing both frequencies would be of great interest."

Why These Technologies Are Kept Secret

The reasons for classification extend beyond the obvious military advantage. The *Electronic Mind Control Book* identifies a deeper motivation: "If everyone had access to such technology, it would quickly destabilize a society. Knowing precisely what an individual is thinking would allow the possibility of significant power over their personality. This means that EM mind control is a sort of 'ultimate weapon' that will probably always be associated either with the unique ability of some esoterics or with extremely secret engineering achievements."

The secrecy serves multiple purposes:

56. **Strategic advantage**: These technologies provide unprecedented capabilities for intelligence gathering, psychological operations, and covert action.
57. **Plausible deniability**: As long as the technologies are classified, governments can deny their existence when victims come forward.
58. **Legal protection**: The absence of acknowledged capability means no legal framework applies, no oversight mechanism exists, and no accountability is possible.
59. **Disinformation**: As Dr. Robert O. Becker noted, "there was a deliberate attempt by authorities to cover up these clandestine activities by means of deliberate disinformation tactics. These included weird and bizarre stories that were leaked to the public for one purpose: to muddy the waters, to make the whole question of mind control seem absolutely unbelievable."

The *Electronic Mind Control Book* concludes that "based on the review of literature made, we deduce which kind of mind reading is scientifically grounded and in what sense we may speak about mind control." The technologies described in this chapter are not science fiction. They are the logical endpoint of research programs that have been funded, classified, and operationalized over decades. The only question is the extent of their deployment — and whether the legal frameworks of the 21st century will be adequate to address them.

CHAPTER 78-bis

Synthetic Telepathy — When It Will Exist, How It Must Be Constrained

Of all the speculative technologies surveyed in Part XV, none commands more public fascination or more public dread than synthetic telepathy: the direct transfer of thought from one brain to another, mediated by technology rather than language or gesture. Chapter 78 mentioned the topic briefly. This chapter treats it at length, because the question of when synthetic telepathy will become possible — and how, if so, it must be constrained — is one of the principal questions the next two decades will pose.

We use the term "synthetic telepathy" with deliberate technical precision. It means: the technological transmission of meaningful semantic content (words, images, concepts, intentions) from the brain of one person to the brain of another, in real time, at a useful rate, without requiring either party to articulate or perceive the content through the ordinary sensory channels of speech, writing, gesture or display. By this definition, the capability does not yet exist, but most of its constituent technologies do.

Synthetic Telepathy — Three Capability Branches and Their Substrates

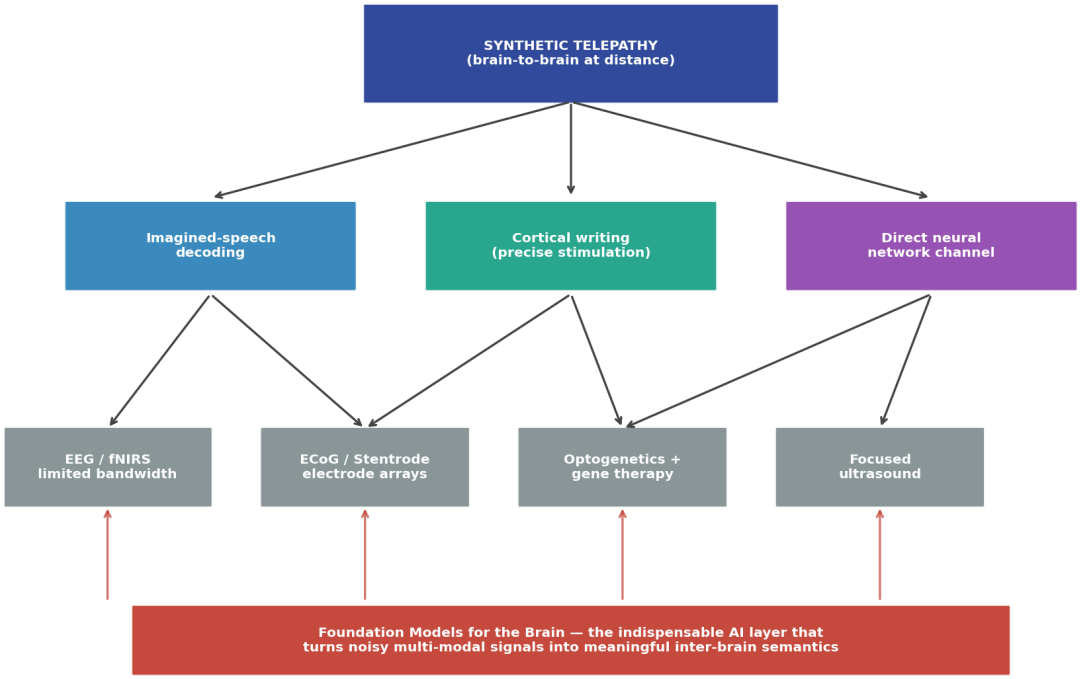


Figure 78bis.1. Synthetic telepathy decomposed into three capability branches — imagined-speech decoding, cortical writing, and direct neural network channels — together with the substrate technologies and the AI inference layer that turns multi-modal signals into inter-brain semantics.

The Three Capability Branches

Branch 1 — Imagined-Speech Decoding

The first branch is reading: capturing the content of a speaker's imagined or attempted speech and rendering it as text or audio. This capability has progressed dramatically in the past five years. Stanford's Krishna Shenoy laboratory and Edward Chang's UCSF group have demonstrated decoding of attempted speech at sixty-plus words per minute in patients with implanted electrocorticography arrays. Brain2Qwerty, GPT-Neuro and related published

systems have demonstrated robust word-level decoding from imagined speech in motor cortex and Wernicke's area at progressively higher accuracy.
The branch's substrate technologies range from low-resolution non-invasive EEG and fNIRS, which produces approximate inferences about coarse cognitive states, through clinical-grade ECoG arrays, which produce word-level decoding of attempted speech, to penetrating electrode arrays such as the Utah Array, which produce the highest-fidelity decoding currently demonstrated. The Synchron endovascular Stentrode represents a middle path: cortical access without craniotomy.
Within the next five to ten years, imagined-speech decoding will plausibly reach the threshold of useful telepathic transmission for select users: paralysed patients whose only realistic communication channel is brain-decoded, surgical patients with temporarily implanted ECoG, and research volunteers. Within ten to twenty years, the combination of foundation models for neural data, higher-resolution non-invasive sensing, and consumer-form-factor brain interfaces could bring crude imagined-speech decoding to mass-market deployment.

Branch 2 — Cortical Writing

The second branch is writing: delivering meaningful content directly to a recipient's brain in a form they perceive as their own thought, sensation or experience. This is harder than reading by several orders of magnitude. The current state of the art is crude: transcranial magnetic stimulation can produce phosphenes (perceived flashes of light); transcranial direct-current stimulation can shift mood and attention; focused ultrasound can stimulate specific subcortical regions; deep-brain stimulation can produce specific sensory and emotional experiences in some patients. None of these is content delivery in the telepathy sense.
The substrate technologies for cortical writing are advancing along three fronts. Optogenetics, which uses light-activated ion channels delivered by viral gene therapy, can in principle deliver high-spatial-resolution writing but requires advance preparation of the target tissue and is currently not approved for human use outside narrow research contexts. Focused ultrasound is non-invasive, increasingly precise, and the most likely route to consumer-grade cortical writing in the medium term. Penetrating electrode arrays, scaled up from current research configurations, can in principle deliver content-rich stimulation but at considerable surgical cost.
Within the next ten to twenty years, cortical writing will plausibly progress from current laboratory demonstrations (perception of simple shapes or sounds delivered directly to the visual or auditory cortex) to delivery of recognisable but limited semantic content (concepts, words). True high-bandwidth content delivery — delivering an entire thought as the speaker conceived it — is plausibly a thirty-to-fifty-year capability if it is achievable at all.

Branch 3 — Direct Neural Network Channel

The third branch is the brain-to-brain channel: a pipeline that captures sender content and delivers it to receiver in some functional form. The first published demonstration of brain-to-brain communication was Rajesh Rao's 2014 collaborative game between two human subjects, in which one subject's imagined hand motion was translated into transcranial magnetic stimulation of the second subject's motor cortex. Since then, multi-subject brain-network demonstrations have been published at small scale.
The channel branch will be the integration of progress in branches one and two: imagined-speech decoding on the sender side, content delivery on the receiver side, and an intermediate AI layer that bridges any semantic mismatch. The principal challenge is bandwidth and fidelity. A high-fidelity channel will require both the sender and receiver to be equipped with high-resolution interfaces, which means cortical implants at present. The consumer-grade telepathy of science fiction — two people in everyday life conversing brain-to-brain through a headband — is, on a realistic technical view, several decades away if it is achievable at all.

The Indispensable AI Layer

Across all three branches, the indispensable common component is the foundation-model layer that turns noisy multi-modal neural signals into semantic content. As Chapter 53 argues, foundation models for the brain represent a qualitative leap in what can be decoded from existing signals. The same models that will make synthetic telepathy possible are the models that will, well before telepathy is realised, make conventional neural reading vastly more invasive. The trajectory of the AI layer is the principal forecasting variable; everything else depends on it.

How Synthetic Telepathy Must Be Constrained

If synthetic telepathy is plausibly achievable in some form within the next two decades, then the question of how it must be constrained is not academic. The following constraints are the minimum that any responsible regulatory regime must establish, and they should be set now, before deployment, so that the technology develops within the constraints rather than against them.

Constraint 1 — Symmetry of Consent

Both the sender and the receiver must consent freely, specifically and revocably to every channel session. The sender's consent is to the transmission of their thought; the receiver's consent is to the reception of foreign content into their cognitive stream. Sender consent without receiver consent is psychic intrusion. Receiver consent without sender consent is theft of thought. Both are required.

Constraint 2 — No Pre-Configuration

It must not be permissible to pre-configure a recipient's neural interface in such a way that future cortical writing can be delivered without per-session consent. This is the analogue of the principle in surgical anaesthesia that a patient's prospective consent does not extend to indefinite future procedures. Each transmission requires fresh consent.

Constraint 3 — Inalienable Content Filter

Certain categories of content — politically coercive material, commercial advertising, content from state actors — must be prohibited from transmission regardless of consent. A person cannot validly consent to being subjected to direct neural advertising any more than they can consent to being enslaved. This is the inalienable-rights principle applied to synthetic telepathy.

Constraint 4 — Mandatory Black-Box Logging

Every channel session must be logged in tamper-evident form by an independent authority, with sender and receiver each receiving a verifiable record of what was transmitted. The recording is required not only for evidence in the event of dispute but also for the recipient's own ability to reflect on, and where necessary to dispute or reject, content delivered into their cognitive stream.

Constraint 5 — Receiver Override

The receiver must be able, at any time, to halt the channel and to expunge from their accessible record any content they wish to dispute or remove. The receiver-override principle is the cognitive analogue of the right to withdraw from a contract or to terminate a medical treatment.

Constraint 6 — Capability Asymmetry Prohibition

Channels in which one party has substantially more capability than the other (higher resolution, broader bandwidth, more advanced AI mediation) must be specifically licensed and audited. The asymmetric channel — in which a powerful sender writes directly to a less-

capable receiver — is the configuration most prone to manipulation and exploitation. It must not become the default.

Constraint 7 — Children and the Cognitively Vulnerable

Synthetic telepathy involving children, persons lacking decision-making capacity, or persons in coercive contexts (prisons, detention, workplaces with severe power asymmetry) must be prohibited absolutely. The capacity for free consent that the technology presupposes is most clearly absent in these populations.

Constraint 8 — Cross-Border Restriction

Channel sessions that cross national borders must be governed by the proposed Global Neurotechnology Authority. The default rule should be that transmission is permitted only between jurisdictions that have ratified the Convention's mutual recognition provisions. This is the only way to prevent the technology being used for cross-border manipulation and surveillance.

The Forum Internum and the Receiver

The deepest constitutional question raised by synthetic telepathy is what it does to the forum internum. The forum internum has historically been protected as absolute because no power could reach it: thought formed within a person's brain remained theirs regardless of external coercion. Synthetic telepathy ends this absoluteness — at least for the receiver. Once content can be delivered directly to a person's cortex, the forum internum is no longer technically inviolate.

The implication is that the absolute prohibition on interference with thought must be reinterpreted not as a prohibition on the technology of transmission (which would foreclose beneficial uses) but as a prohibition on unauthorised reception. The receiver-consent and receiver-override principles above are the operational expression of this reinterpretation. A person retains sovereignty over the cognitive stream that constitutes their inner life; what they choose to let in is theirs to choose, and what is delivered without their choice is a violation regardless of the sender's intent.

When Does It Become Real?

The honest scientific answer is: it depends on which branch and at what fidelity. Crude imagined-speech decoding in clinical patients is real now. Useful imagined-speech decoding in research volunteers will be real within five years. Mass-market consumer-grade imagined-speech decoding from non-invasive interfaces is plausibly a ten-year prospect. Cortical writing capable of delivering recognisable semantic content is plausibly fifteen years out. High-bandwidth two-way channels between cognitively healthy users at consumer cost are plausibly thirty years out, if achievable.

The implication for policy is that the question is not whether synthetic telepathy will require regulation but when. The constraints above should be drafted now, embedded in the proposed UN Neurotechnology Convention, and elaborated in national statutes, so that as the capability matures the regulatory frame is already in place. This is the rare case where law has the chance to precede deployment rather than chase it. We should take the chance.

PART XV-bis

Hypothetical Cases — Fictional Stories for Reflection

The chapters in this part are works of fiction. None of the persons, companies, devices, agencies or events described is real. The names are invented. Any resemblance to real cases is coincidental except where a real case is explicitly cited as background. The stories are offered as food for thought — narrative devices for thinking through how existing law would apply, where it falls short, and what kind of jurisprudence or legislation would be needed to close the gap.

Each story is told briefly, with sufficient detail to allow the reader to identify the legal issues and to follow the analysis. After each story we offer a short legal commentary, distinguishing

what existing law can already do from what new doctrine or new legislation would be required. The cases are deliberately spread across the next decade — set between 2026 and 2036 — so that the reader can see the trajectory along which the legal challenge unfolds.

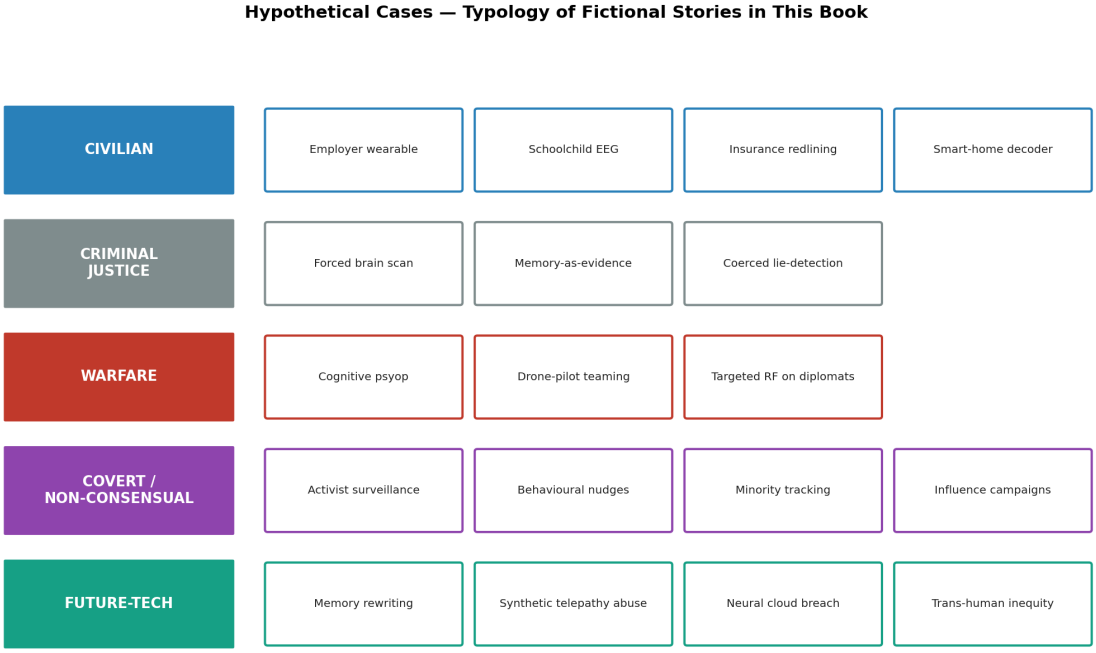


Figure XVbis.1. Typology of the fictional stories in this part. Civilian, criminal-justice, warfare, covert, and future-technology categories together cover the principal classes of dispute that will arise over the next decade.

CHAPTER XV-bis-1

Civilian Case One — The Employer Wearable (2027)

[This is a fictional story. The persons, company, and events described are imaginary. The legal analysis is offered for instructive purposes.]

Marta works as a customer-service representative at a mid-sized Belgian insurance company. In January 2027 the company introduces a new employee-wellness programme. Every employee receives a Muse-style EEG headband. Use is technically voluntary — the company emails a forty-page consent form and asks employees to sign — but the bonus pool for the quarter is restructured so that participants receive a measurable benefit. Marta signs.

Six months in, Marta's annual review is unusual. Her supervisor refers to "attention metrics" Marta has never seen. Marta asks for the underlying data and is told that the metrics are proprietary, that she signed away her right to the raw data, and that she should focus on her performance rather than the methods. The metrics determine ten per cent of her bonus. Marta refuses to wear the device for the next quarter. Her bonus is cut by thirty per cent. Marta consults a lawyer. Three legal frames are available to her in 2027.

First, the EU AI Act's Article 5 prohibits the placing on the market, putting into service, or use of AI systems intended for inferring emotions in the workplace and educational settings (with narrow exceptions). The metrics derived from Marta's EEG arguably constitute emotion inference. The AI Act's mechanism for enforcement is administrative: the Belgian AI Office can act on a complaint, ultimately fining the company a percentage of global turnover.

Second, the General Data Protection Regulation treats biometric data as a special category. Marta's consent is suspect under Article 7(4) because of the imbalance of power between employer and employee; the Article 29 Working Party has long argued that employee

consent is rarely free. The Belgian Data Protection Authority can investigate, order data deletion (echoing the Chilean Emotiv remedy), and issue a fine.
Third, Belgian labour law's prohibition on certain forms of workplace surveillance applies to the device. The labour inspectorate can investigate, particularly because the inferences were used for performance scoring without disclosure of the methodology to the affected employee.
What existing law can already do: Marta probably wins. The three frames operate together. The AI Act provides the cleanest path; the GDPR provides the data-deletion remedy; labour law provides the back pay. Marta's case is appealable through standard administrative and judicial channels. The remedies are reasonable but not transformative.
Where existing law falls short: even if Marta wins, three problems remain unaddressed. First, the underlying data — six months of Marta's neural recordings, plus the inferences derived from them — has already been generated, processed, possibly transferred to cloud-service providers and possibly used to train models that will continue to be deployed. The deletion order may not reach all of it. Second, the case takes two to three years to resolve. Marta has needed the bonus the whole time. Third, the case sets only a narrow precedent; the next case at the next company starts over.
Where new jurisprudence or legislation is required: the ideal regime, developed in Chapter 67-bis, would treat workplace neural monitoring as one of the inalienable cores — prohibited regardless of consent. The Belgian Constitution does not yet contain a neuro-rights clause; a Chilean-style amendment would have foreclosed the underlying programme. The European Union should explicitly add neural inference to the AI Act's per-se prohibited categories in workplace contexts, removing the current exemption pathway. And the Belgian Code of Civil Procedure should be amended to allow representative actions in neuro-rights cases, so that Marta's case sets a precedent for her colleagues without requiring each of them to litigate separately.

CHAPTER XV-bis-2

Civilian Case Two — The Schoolchild EEG (2028)

[This is a fictional story. The persons, school, and events described are imaginary.]
Lucas is ten years old. His school in São Paulo introduces an "attention enhancement" programme in third grade: every child wears a low-cost EEG headband during three afternoon lessons a week. The device displays a colour-coded attention score on the teacher's screen and feeds data to a Brazilian start-up that has partnered with the school district. Parents receive a consent form at the start of the year. Lucas's parents sign because the alternative — opting their child out — would mean Lucas missing the lessons or being placed in a separate group during them.
After six months, Lucas — who is dyslexic and who has previously had no behavioural difficulties at school — is described in his term report as having "sustained attention anomalies" relative to his peer baseline. The school recommends a referral to a paediatric neurologist. The neurologist, examining Lucas, finds nothing unusual. Lucas's parents discover by chance that the attention data has been incorporated into Lucas's permanent school record, that the start-up retains it for ten years, and that Lucas's data has been shared with two universities for research purposes. The consent form mentioned none of this in language a parent could reasonably understand.
Lucas's parents consult a lawyer. The legal frames available to them in 2028 include the Brazilian Lei Geral de Proteção de Dados (LGPD), the Brazilian Constitution's amendments on mental integrity (passed in 2024, modelled on Chile), the Brazilian Statute of the Child and Adolescent, and the UNESCO Recommendation on the Ethics of Neurotechnology as soft-law interpretive guidance.
What existing law can already do: a great deal. The LGPD treats children's data as a heightened category requiring specific, informed parental consent; the consent obtained does not plausibly satisfy this standard. The mental-integrity amendment supplies the

constitutional backdrop and allows the court to read the LGPD's requirements stringently. The Statute of the Child and Adolescent allows the Public Ministry to intervene on Lucas's behalf. The Public Ministry of São Paulo investigates; the school district suspends the programme; the start-up is fined; Lucas's data is deleted.

Where existing law falls short: Lucas's parents' victory does not address several deeper issues. The inference about Lucas — that his attention pattern is anomalous — has already entered Lucas's school file and influenced how his teachers think about him. Even with deletion, the effect of the original characterisation may persist. The two universities that received the data may have already trained models on it; the models continue to exist even after the data is deleted. Other schools in Brazil continue similar programmes; without national legislation, the parents in each must litigate separately.

Where new jurisprudence or legislation is required: a national prohibition on neural monitoring in primary and secondary schools, modelled on the EU AI Act's narrow workplace ban but extended to education, would foreclose the underlying programme. The Brazilian Federal Council on Children's Rights should issue a binding standard prohibiting neural monitoring of minors. The Brazilian Ministry of Education should withdraw approval of any curriculum that incorporates neural monitoring.

The deeper question Lucas's case poses is whether children can ever validly consent — through their parents — to neural monitoring that will follow them through their educational lives. The argument developed in Chapter 59 is that they cannot, and Lucas's case is one of many that will press the question into Brazilian and other jurisprudence over the coming years.

CHAPTER XV-bis-3

Civilian Case Three — Insurance Redlining (2030)

[This is a fictional story.]

Aïcha applies for a fifteen-year mortgage from a mid-sized French bank in early 2030. The application is processed through an algorithmic underwriting system. Her application is approved at a rate 0.6 percentage points above the published median for her income and credit-score profile. Aïcha asks why.

The bank, citing trade-secret protections in its underwriting model, declines to explain. Aïcha files a complaint with the French Commission Nationale de l'Informatique et des Libertés (CNIL). Through the investigation it emerges that the model incorporates, among hundreds of features, a credit-score adjustment derived from sleep-pattern variability in data licensed from a major consumer-wearable company. Aïcha's sleep pattern, measured by a fitness tracker she has worn for three years, shows variability classified as a risk factor for long-term solvency.

Aïcha had no idea her sleep data was being used this way. The wearable's terms of service had granted the company broad rights, including the right to share with financial-services partners. Aïcha had not read the terms. Few people do.

The legal frames available in 2030 include the GDPR, the EU AI Act (now somewhat extended by post-2026 amendments), the French Consumer Code, the AI Act's specific provisions on credit-scoring (high-risk under Annex III), and the early French national neuro-rights statute (passed in 2028 in response to the EU regulation).

What existing law can already do: CNIL can investigate. The high-risk classification of the AI system means it must comply with extensive transparency and bias-audit requirements. The Article 22 GDPR right not to be subject to a decision based solely on automated processing applies. The Consumer Code requires meaningful explanation of credit decisions. Aïcha is likely to win a finding of violation, a re-underwriting of her application without the sleep-pattern feature, and damages.

Where existing law falls short: Aïcha's case reveals a structural problem that her individual remedy does not solve. The bank's model has been trained on millions of borrowers, many of whom have similar sleep-pattern characteristics and have been similarly disadvantaged

without ever filing a complaint. The original transfer of sleep data from the consumer wearable to the bank — through a chain of data brokers — may not have been illegal at the time it occurred. The inference itself — that sleep variability predicts solvency — may or may not be valid; even if statistically valid, it may be discriminatory if it correlates with protected characteristics.

Where new jurisprudence or legislation is required: a categorical prohibition on the use of any neural-derived inference in credit and insurance underwriting, modelled on the existing prohibition on genetic-information use in many jurisdictions. The French national statute should be amended to make this explicit. The EU AI Act should be amended to add neural inferences to its high-risk Annex III with a specific prohibition for credit and insurance. And the practice should be made an inalienable prohibition: not waivable even with consent, because the consent obtained at the wearable's signup is structurally inadequate.

CHAPTER XV-bis-4

Civilian Case Four — The Smart-Home Decoder (2031)

[This is a fictional story.]

A US-based home-electronics company introduces a premium smart-speaker system that includes a near-field EEG sensor in the bedside variant. The sensor, marketed as a sleep-quality improvement feature, gathers low-resolution EEG during sleep. The company claims the data is processed entirely on-device and never leaves the home. In 2031, a security researcher publishes evidence that the encrypted device-to-cloud channel includes the EEG data; the company corrects the marketing language but continues the underlying data-collection.

Among the device's users is Eleanor, a retired Federal judge in California. Eleanor's husband had bought the speaker as a Christmas gift. When the security researcher's report appears, Eleanor is concerned and asks the company for her data; she receives, in encrypted format, eighteen months of nightly EEG recordings.

Eleanor consults former colleagues, who consult an expert in the (rapidly-developing) doctrine surrounding the California Consumer Privacy Act, as amended in 2024 to cover neural data, and the federal MIND Act, passed in 2030. The CCPA amendment specifically requires opt-in for sensitive neural data; the device's terms of service offered opt-out only. The MIND Act, as enacted, prohibits the sale of neural data without explicit consent and creates a private right of action.

What existing law can already do: a great deal. Eleanor's case is well-positioned. The CCPA violation supports state-level remedies (statutory damages of $750 per violation per consumer, multiplied by the consumer class). The MIND Act violation supports federal remedies. The class is large — millions of devices — and the damages are substantial.

Where existing law falls short: Eleanor's case will produce a settlement; the company will pay several billion dollars; the law firms will prosper. But the underlying capability — the ability of consumer-electronics manufacturers to embed neural sensing in everyday devices without users' clear understanding of what is being measured and for what purpose — will not be addressed. The remedy is monetary; the structural change is not.

Where new jurisprudence or legislation is required: a categorical labelling requirement, like the FCC equipment-registration label but for neurotech: every device that can sense neural activity must carry a standardised label visible at purchase, with a national registry of such devices. The MIND Act should be amended to require this. The FCC should be designated as the implementing agency. And the black-box logging requirement developed in Chapter 67-bis should apply to consumer devices as well as clinical ones: every neural sensor must maintain a tamper-evident log of what it has measured and where the data has flowed.

CHAPTER XV-bis-5

Warfare Case One — Cognitive Psyop on Civilians (2028)

[This is a fictional story. The conflict described is imaginary.]
A regional conflict in 2028 between two mid-sized states includes, on one side, a sustained cognitive-warfare campaign against the other side's civilian population. The campaign is conducted through a combination of foreign-language social-media amplification, deepfake video and audio targeting specific community leaders, microtargeted advertising informed by a deep dataset of psychological profiles (some of which derive ultimately from consumer neurotech wearables marketed years earlier under benign labels), and synchronised attacks on the cognitive resilience of the target population over a six-month period.
Documented effects include a measurable shift in civic trust, declining electoral participation, the breakdown of community-level conflict-resolution norms, and a substantial increase in inter-community tension. The attacking state's deniability is plausible at every individual operation but the cumulative pattern is unambiguous.
The target state's response is constrained. The Geneva Conventions do not address cognitive warfare directly; the protections for civilians under the Fourth Geneva Convention apply but their evidentiary thresholds are designed for kinetic harms. The 1899 Hague Convention's prohibition on poisoned weapons does not extend by analogy. The Tallinn Manual on the International Law Applicable to Cyber Operations provides interpretive guidance but is not binding and does not address the specific question of cognitive harm. NATO's cognitive-domain doctrine acknowledges the threat but provides primarily defensive guidance.
What existing law can already do: little, with specificity. International humanitarian law's general protections for civilians apply at a level of abstraction that does not bind the attacking state to specific obligations. The International Criminal Court could in principle prosecute attacks that meet the threshold of crimes against humanity, but the evidentiary challenge is enormous: how does one prove that an electoral participation decline was caused by a cognitive-warfare operation rather than ordinary political dynamics?
The customary international law of countermeasures may justify some retaliatory action by the target state, but the absence of an agreed definition of what counts as cognitive-warfare aggression makes proportionality assessment difficult and risks escalation.
Where new jurisprudence or legislation is required: a binding international convention on cognitive warfare, developed under the auspices of the proposed UN Neurotechnology Convention. The convention should: define cognitive-warfare aggression with specific reference to mass-scale, intentional, technologically-mediated impairment of civic deliberation in another state; prohibit such operations in peacetime under the same logic as the prohibition on aggressive war; establish a verification regime that can produce attributable findings; create a mechanism for collective response analogous to the UN Security Council's Chapter VII authority. The challenge is that the attacking states in actual cases are likely to be permanent members of the Security Council; the convention must therefore include alternative response mechanisms.

CHAPTER XV-bis-6

Warfare Case Two — Drone-Pilot Brain-Computer Teaming (2032)

[This is a fictional story.]
By 2032, the air forces of three major powers have deployed brain-computer-interface augmentation for drone pilots. The interface, a high-channel-count non-invasive system, allows the pilot to direct multiple drone systems simultaneously by intention; closed-loop feedback adjusts the pilot's attention and stress state to optimise performance. The systems

have been deployed in a regional conflict in which the augmenting state's drones have engaged forces of a non-augmenting state.
Operator A, a drone pilot in the augmenting state's air force, conducts a strike that kills a number of non-combatants. The post-incident review establishes that the closed-loop feedback system was actively adjusting Operator A's stress and attention during the engagement; specifically, the system's algorithm reduced Operator A's perception of moral hesitation, on the grounds that hesitation was correlated with poor mission outcomes in training data. Operator A asserts at the inquiry that they would not have conducted the strike absent the augmentation.
The legal frames available include international humanitarian law (the principles of distinction, proportionality, and precaution), the Rome Statute's individual-responsibility provisions, the augmenting state's national military justice code, and an emerging body of doctrine on autonomous and augmented systems.
What existing law can already do: the strike is reviewed for IHL compliance regardless of augmentation. If the strike was disproportionate or failed to take feasible precautions to spare civilians, the responsible commanders can be prosecuted under the Rome Statute. Operator A's individual liability is more complex: the augmenting state's command structure may share liability if the system was foreseeably likely to produce such outcomes.
Where existing law falls short: the augmentation introduces a novel question of cognitive autonomy at the operational level. If Operator A's perception of moral hesitation was technologically suppressed, the question of whether Operator A's decision was meaningfully Operator A's becomes acute. International humanitarian law presumes a moral agent making decisions; the augmentation arguably impaired that agency. But Operator A consented to the augmentation upon enlistment, and the closed-loop system was authorised by competent command. The case sits in a doctrinal gap.
Where new jurisprudence or legislation is required: international humanitarian law should be supplemented by an agreed standard for the cognitive autonomy of military operators in augmented contexts. The standard should distinguish: (a) augmentation that enhances baseline human capabilities while preserving the operator's full moral agency, which is broadly acceptable; (b) augmentation that selectively suppresses categories of moral reasoning (hesitation, empathy, etc.) in service of operational effectiveness, which should be prohibited as inconsistent with IHL's presumption of moral agency; (c) closed-loop systems that substitute algorithmic decisions for human ones, which fall under the existing emerging consensus on lethal autonomous weapons systems. National military justice codes should be amended to recognise the third category as a per-se violation. The augmenting state should be required to maintain detailed logs of closed-loop system interventions and to make them available to post-incident review.

CHAPTER XV-bis-7

Warfare Case Three — Targeted RF on Diplomats (2029)

[This is a fictional story.]
In 2029, eleven diplomats from a Western state stationed at the embassy of a host state in a third state report a cluster of symptoms consistent with the Havana Syndrome profile described in Chapter 26: sudden onset of pressure, pain, dizziness, cognitive fog. Two of the diplomats experience lasting cognitive impairment. The dispatching state, working from an updated national-academies report on the now-better-established science of pulsed radiofrequency exposure, concludes that the cluster is more likely than not the result of intentional targeting. Attribution is uncertain between the host state and a third-party actor with the capability to deploy in the host state's territory.
The dispatching state's options are constrained by the same limitations that constrained the original Havana cases. There is no binding international instrument prohibiting the

development or use of neuroweapons. The Vienna Convention on Diplomatic Relations protects diplomats from various forms of interference but does not specifically address neuroweapons. The Geneva Conventions apply only in armed conflict, which is not the situation. The Convention against Torture applies to states-as-actors against persons in their custody or jurisdiction.
What existing law can already do: the dispatching state can declare host-state diplomats persona non grata; the dispatching state can suspend diplomatic relations; the dispatching state can take economic countermeasures; the dispatching state can bring the matter to the UN Security Council, where its effectiveness depends on the constellation of permanent members' interests. None of these responses produces a finding of legal violation by a specific actor; none deters future operations against other targets.
Where existing law falls short: the absence of a binding prohibition on neuroweapons, of a verification regime, and of an internationally-recognised attribution mechanism makes any individual incident effectively unsolvable as a legal matter. The dispatching state's options reduce to diplomatic and political tools that do not constrain the underlying capability.
Where new jurisprudence or legislation is required: this is the case most clearly demanding the proposed UN Neurotechnology Convention developed in Chapter 68 and Chapter 67-bis. The Convention's specific provisions on neuroweapons should: prohibit the development, production, stockpiling, transfer or use of devices intended to covertly impair or control cognition; establish a verification regime modelled on the OPCW; create an attribution-and-response mechanism that can act on incidents like this one; and provide for individual criminal liability for senior officials who order or facilitate such operations. The dispatching state's individual response options would remain available as countermeasures within the Convention framework, but the Convention would supply the legal anchor that the dispatching state's response currently lacks.

CHAPTER XV-bis-8

Covert Case One — Activist Surveillance (2027)

[This is a fictional story.]
Anna is a journalist in an authoritarian state covering, with international support, the activities of a particular state-security service. In 2027, Anna receives an anonymous package containing what appears to be a consumer-grade wireless EEG headband, originally manufactured in a different country, modified with what a friendly technical adviser identifies as a custom firmware that streams raw EEG data over the local cellular network using a steganographic protocol. The headband had been provided to Anna six months earlier by a source she had trusted; the source has since become unreachable.
Anna had worn the headband — believing it to be a sleep-quality tracker — during a number of sensitive interviews. The data, if it reached the state-security service, would constitute extensive neural-derived intelligence about Anna's mental state during specific moments of journalistic activity, the identities of her sources (inferred from her stress responses during identifications), and her physical movements over a six-month period.
Anna is now in a position familiar to victims of covert surveillance: she has reason to believe she has been compromised but cannot easily prove it, the state has plausible deniability, and the international legal frame is limited.
What existing law can already do: the manufacturer's terms of service have been technically violated (the firmware modification), giving the manufacturer cause of action; this is of cold comfort to Anna. The general international human-rights instruments protect Anna's freedom of expression (ICCPR Article 19) and privacy (Article 17); their enforcement depends on the state in question's submission to international mechanisms. International press-freedom NGOs can document the case and use it for advocacy.
Where existing law falls short: the case demonstrates the structural inadequacy of the consent-based neuro-rights regime when the consent has been obtained by deception. Anna consented (implicitly, by wearing the device) but her consent was vitiated by

misrepresentation. There is no specialised forum for adjudication; the state-security service's incentive to deny is total.
Where new jurisprudence or legislation is required: a specific international crime of covert neural surveillance, with universal jurisdiction in the convention sense, would allow Anna's case to be prosecuted by any state party that gained jurisdiction over a responsible individual. The proposed Convention should include such a provision. National criminal codes should be amended to create the corresponding domestic offence. The wearable manufacturer should be required, as a condition of placing devices on the market, to maintain a cryptographic chain-of-trust on its firmware that would make this kind of modification impossible without detection at the device's next contact with the manufacturer's servers.
The case also illustrates the importance of detection infrastructure (Chapter 40 and Chapter 71). A regulatory regime that required consumer neurotech devices to attest their firmware integrity to an independent authority on a regular basis would have allowed Anna to know within hours that her device had been tampered with.

CHAPTER XV-bis-9

Covert Case Two — The Untraceable Behavioural Nudge (2030)

[This is a fictional story.]
In 2030, a privately-held American company markets a consumer-grade sleep-improvement device that delivers gentle transcranial direct-current stimulation during the night. The device is marketed honestly: users are told the device delivers small electrical currents to the scalp to support sleep quality. Users sign a clear consent form. The device receives FDA over-the-counter approval as a wellness device.
Internally, the company's product team — without the knowledge of the company's senior leadership or the FDA — modifies the stimulation profile to include subtle parameters intended to increase the user's general suggestibility on a measurable scale. The product team's stated goal is to make the device's advertising and recommendation features more effective, increasing both user retention and revenue from partner companies.
The modification is discovered three years later, when a security researcher publishes a comparative analysis of the device's stimulation profile against published norms. By that time, approximately 1.4 million devices have been deployed. The researcher's findings are confirmed by independent academic analysis. The company denies that the modified parameters had the intended effect (the relevant evidence is published and contested).
The legal frames available include the federal Food, Drug, and Cosmetic Act (the modification arguably converts a wellness device into an unapproved medical device, which is a federal crime), the FTC Act (deceptive trade practices), state consumer-protection statutes, and the new federal MIND Act.
What existing law can already do: a great deal of enforcement is possible. The FDA can pursue criminal prosecution of responsible officers under the Park doctrine, which extends liability to corporate officers who could have prevented the violation. The FTC can pursue deceptive-practices remedies including disgorgement of profits. State attorneys general can pursue parallel actions. Class-action litigation by affected consumers is virtually certain to follow.
Where existing law falls short: the substantive question of whether the modified parameters actually produced behavioural change in users is contested and may take years to resolve. Even if resolved against the company, the question of individual harm is acute: 1.4 million users have been exposed for years; some have made significant decisions during the period (purchases, marriages, employment changes) that they cannot easily unmake. The harm is real and unrecoverable, even if the legal remedies are pursued.

Where new jurisprudence or legislation is required: a categorical prohibition on cortical writing that is intended to influence behaviour beyond the device's disclosed therapeutic purpose, with strict liability and substantial criminal penalties. This should be one of the inalienable cores of the proposed UN Convention. The MIND Act should be amended to make this prohibition explicit at federal level. The black-box logging requirement should apply: every stimulation device must log every stimulation event, the parameters used, and the firmware version, in tamper-evident form available to FDA inspectors and to affected consumers in the event of dispute.
The case also makes the structural argument for inalienable rights: even if the company's users had explicitly consented to the modified parameters, the consent should not have been honoured. There are interventions on the human cognitive apparatus that should be prohibited regardless of consent because they bypass the conditions for valid consent.

CHAPTER XV-bis-10

Covert Case Three — Minority Tracking (2031)

[This is a fictional story.]
In 2031, an investigation by a consortium of investigative journalists uncovers that the security services of a particular state have, over a five-year period, used a combination of consumer-wearable data (purchased through data brokers), public-camera face recognition, and inference models trained on public-broadcast neuroscience research to develop a profile of physiological markers associated with members of a particular religious minority. The profile is used to flag individuals at airport security and at internal checkpoints for additional scrutiny.
The investigation reveals approximately 14,000 individuals flagged across the five-year period, of whom 280 were detained for periods ranging from hours to weeks, of whom 12 were prosecuted under counter-terrorism law on charges that the investigation concludes were largely or entirely based on the flagging system.
The legal frames available include the state's national constitutional protections against religious discrimination (varying by jurisdiction), the International Covenant on Civil and Political Rights' protections against arbitrary detention (Article 9) and discrimination (Article 26), the European Convention on Human Rights (if the state is a party), and the proposed UN Convention's anti-discrimination provisions.
What existing law can already do: international and regional human-rights bodies can take cognisance of the practice. The European Court of Human Rights, if the state is a party, can issue binding judgments. Individual victims can bring domestic and international litigation. The international press attention is itself a form of pressure.
Where existing law falls short: the underlying technological practice — the use of physiological inferences to flag minorities — falls into a doctrinal gap. The practice is discriminatory in effect but its discriminatory intent is officially denied; the state defends the practice on neutral security grounds. The inference models can be cited as proprietary; the procurement chain (consumer data → brokers → state) is opaque; the individuals affected lack standing in many domestic forums; the cumulative effect is significant but resists single-case remedy.
Where new jurisprudence or legislation is required: a specific prohibition on the use of neural and physiological inferences in security profiling, with substantive accountability for the procurement chain. The proposed Convention should include such a prohibition with extraterritorial effect: a state that uses such inferences against the citizens of another state party should face Convention sanctions. National law should require warrants for any use of neural-derived inferences in security operations, with judicial oversight.
The case illustrates the cross-border dimension of the problem: the consumer-wearable data was collected lawfully in other states; the brokers were international; the inference models were trained on globally-available public research. National law alone cannot solve a fully-international problem.

CHAPTER XV-bis-11

Covert Case Four — Foreign Influence Campaign (2032)

[This is a fictional story.]
Two months before a parliamentary election in a mid-sized European state, a foreign-state-affiliated entity conducts what intelligence services later identify as a sophisticated cognitive-warfare operation. The operation does not rely on visible disinformation; instead, it uses microtargeted advertising on three major platforms, informed by a deep dataset purchased through legitimate-appearing data brokers, to subtly shift attentional patterns and emotional valences of approximately 800,000 swing voters in particular constituencies.
The shift is sub-conscious — voters do not perceive themselves as having been influenced — but post-election analysis suggests a measurable effect on the outcome in seven constituencies, sufficient to alter the composition of the parliamentary majority by a small but consequential margin.
The legal frames available include the state's electoral law (which addresses overt foreign interference but is ambiguous on this kind of operation), the EU Digital Services Act (which addresses platform conduct), the EU AI Act (which prohibits certain manipulative AI techniques), the General Data Protection Regulation (which limits the underlying data flows), and the Council of Europe Convention on Cybercrime.
What existing law can already do: significant enforcement against the platforms is available under the Digital Services Act and the AI Act. The data brokers face GDPR liability. The state's election commission can investigate. The foreign-state-affiliated entity itself, however, is largely beyond reach of the state's domestic law.
Where existing law falls short: the operation falls in a category that election law was not designed for. The voters were not deceived in the conventional sense; they were nudged by techniques that bypassed their conscious deliberation. The platforms argue, with some force, that they did not knowingly facilitate the operation. The foreign-state-affiliated entity denies the operation. The damage to the legitimacy of the electoral outcome is real but resists conventional remedy: a re-run of the election would not undo the cognitive effects, and there is no constitutional mechanism in the affected state for retrospective adjustment.
Where new jurisprudence or legislation is required: the proposed Convention should treat cognitive-warfare operations against electoral processes as a category subject to specific prohibition and collective response. National electoral law should be amended to address sub-conscious manipulation explicitly. Platform-level rules should require pre-election cooling-off periods on microtargeted political advertising. And there must be a normative red line: deliberate cognitive operations against an election are an act of aggression analogous to (and possibly more consequential than) the conventional armed-conflict thresholds the United Nations Charter contemplates.

CHAPTER XV-bis-12

Future-Tech Case One — The Memory Rewrite Trial (2034)

[This is a fictional story based on a fictional future breakthrough.]
By 2034, a clinical-research consortium has demonstrated, in a small but well-controlled trial, that a specific combination of pharmacological agents and targeted optogenetic stimulation can selectively weaken the consolidation of traumatic memories. The trial has reported significant clinical benefit for veterans with treatment-resistant post-traumatic stress disorder, with no apparent side effects in the eighteen-month follow-up window.

A small private clinic in a permissive jurisdiction announces that it will offer the procedure on a fee-for-service basis to non-veterans, expanding the indication to include traumatic memories from civilian contexts (assault, accidents, abuse). Clients include the survivors of high-profile assault cases who are scheduled to testify in upcoming prosecutions.
A defendant in one of these prosecutions discovers that the principal witness against him has undergone the procedure to reduce the trauma associated with her testimony. The defendant moves to exclude her testimony on the grounds that the witness has been technologically altered in ways that affect the reliability of her memory.
The legal frames available include the rules of evidence in the relevant jurisdiction (which require witness testimony to reflect the witness's genuine recollection), the constitutional right of the defendant to confront and effectively cross-examine, the witness's right to medical treatment, and the proposed UN Convention's emerging provisions on consent and reversibility.
What existing law can already do: the court must adjudicate the motion. Rules of evidence will be applied; the court will consider whether the witness's altered memory satisfies the standards for reliable testimony. Expert testimony from neuroscientists will be heard on both sides. The judge will make a ruling that may be appealed.
Where existing law falls short: the underlying scientific question — whether a memory altered by the procedure remains the same memory in the relevant legal sense — has no established answer. The rules of evidence presuppose a kind of memory continuity that the procedure has, by design, disrupted. The witness has a strong interest in medical treatment; the defendant has a strong interest in reliable evidence; these interests collide directly.
Where new jurisprudence or legislation is required: a specific evidentiary doctrine for technologically-altered memories. The rules of evidence should require: (a) full disclosure to the court of any neurotechnological intervention affecting the witness's memory; (b) expert evaluation of the specific effects on the testimony in question; (c) consideration of the timing of the intervention relative to the events and the testimony; (d) appropriate cautionary instructions to the trier of fact.
The case also raises a deeper normative question that the UN Convention should address: should there be a categorical prohibition on memory modification of witnesses in pending proceedings? The argument for prohibition is the integrity of the evidentiary process; the argument against is the witness's right to medical treatment. The Convention's position, in the model legislation of Chapter 67-bis, is that the witness's right to treatment prevails, but only with full disclosure to the proceedings and the application of structured evidentiary safeguards.

CHAPTER XV-bis-13

Future-Tech Case Two — The Synthetic-Telepathy Abuse (2036)

[This is a fictional story based on a fictional future breakthrough.]
By 2036, two large technology platforms have launched competing consumer-grade synthetic-telepathy services, on the constraints elaborated in Chapter 78-bis. The services allow consenting paired users to exchange brief semantic content through high-resolution non-invasive interfaces. The services are marketed for intimate communication between partners, for collaborative creative work, and for accessibility uses. They are subject to the Convention's emerging constraints, including symmetric consent, mandatory black-box logging, and receiver override.
Diego and Lara, a couple in Buenos Aires, are paired users of one of the platforms. During an emotionally heated argument, Lara — who is the sender at the moment — transmits a sustained burst of content that the platform's content-filter identifies as crossing the threshold of psychologically abusive material. The platform's automatic safeguards stop the transmission; Lara is briefly suspended from the platform under its terms-of-service rules.

Diego files a complaint with the Argentine consumer-protection authority and considers civil action for emotional damages. Lara argues that the platform's content filter is overbroad, that her transmission was within the scope of normal interpersonal communication, and that the platform's intervention has itself caused her harm by characterising her as abusive in the platform's records.
The legal frames available include the Argentine Civil Code's general provisions on harm and consent, the Argentine constitutional protections for personal communication, the emerging Argentine neuro-rights statute (passed in 2032), the platform's terms of service, and the proposed UN Convention's provisions on synthetic telepathy.
What existing law can already do: civil-action remedies are available to Diego; administrative-complaint remedies are available to Lara. The platform's interventions can be scrutinised under consumer-protection law. The constitutional protections for personal communication apply, with the novel question being whether they apply differently to direct-neural transmission than to conventional speech.
Where existing law falls short: the substantive question of what content is permissible in synthetic telepathy is new and contested. The platform's content-filter is functioning as a private speech regulator over the most intimate possible communication channel. The classification of Lara's transmission as abusive has been made by an AI system whose reasoning is opaque. The dispute resolution mechanisms have not been designed for this kind of dispute.
Where new jurisprudence or legislation is required: the proposed UN Convention's synthetic-telepathy provisions, and the national implementing legislation, must address several specific issues. First, the standard for permissible content in private synthetic-telepathy transmissions should be set by the receiver's revocable consent rather than by platform judgement; the platform's role should be to enforce the receiver's preferences, not to substitute its own. Second, the receiver-override principle should give Diego the ability to halt the transmission and to expunge any content he wishes; this is the appropriate location of the protection. Third, the content-filter classification of Lara's transmission as abusive should not be incorporated into Lara's permanent record without due process; the platform must provide a meaningful right of review.
The case illustrates the structural challenge that platform governance poses for synthetic telepathy. The platform's power to mediate the most intimate human communication is enormous; its current accountability is the same as any social-media platform's. The Convention's provisions must include specific platform-governance rules for synthetic telepathy that exceed those for conventional communication: independent ombudsperson, transparency on filter operation, due process for users.

CHAPTER XV-bis-14

Future-Tech Case Three — The Neural Cloud Breach (2033)

[This is a fictional story based on fictional future technology.]
By 2033, an estimated forty thousand patients worldwide have received cortical implants. The major implant manufacturers maintain cloud-based data services for clinical management, firmware updates, and analytic feedback. The manufacturers have varying security postures.
In late 2033, a security researcher discloses to one of the manufacturers a critical vulnerability in their cloud services that would allow an attacker to read recorded neural data from all implanted devices and, in theory, to issue stimulation commands. The vulnerability is being responsibly disclosed; the manufacturer has ninety days to patch it. Before the patch is deployed, however, the vulnerability is independently discovered and exploited by an adversary with state-level capabilities. The adversary reads — and exfiltrates — six months of

stored neural data on approximately twelve thousand patients, including identifiable medical and quotidian content.
The discovery of the breach is made within days; the adversary's identity is plausibly attributable to a state-affiliated unit; the patients are notified within weeks. The full extent of the data exposure is not immediately knowable because the manufacturer's logging was inadequate.
The legal frames available include national data-protection statutes (LGPD, GDPR, MIND Act, etc.), national medical-device regulation, breach-notification statutes, and the proposed UN Convention's provisions on cross-border neural data flows.
What existing law can already do: substantial regulatory action against the manufacturer is available. Data-protection authorities can impose significant fines; medical-device regulators can issue corrective-action orders; affected patients can pursue civil litigation. The aggregate financial exposure of the manufacturer is in the multi-billion-dollar range.
Where existing law falls short: the substantive harm is largely unrecoverable. Patients' neural recordings have been exfiltrated; the adversary's copies cannot be retrieved. The patients face indefinite uncertainty about what use will be made of their data. Some patients will respond by requesting explantation of their devices, which is a major surgical intervention with substantial risks. The trust placed in the underlying technology is shaken across the user community.
Where new jurisprudence or legislation is required: a specific neural-data security regime substantially more stringent than general data protection. The proposed Convention should require: (a) substantive security standards for systems handling neural data, modelled on the FAA's airworthiness standards for aircraft systems rather than on general data-protection guidance; (b) mandatory black-box logging that would have made the extent of the breach knowable immediately; (c) substantive consequences for breach (not just financial liability but mandatory product recalls and ongoing regulatory oversight); (d) treatment of state-actor exfiltration of neural data as an act of aggression triggering the Convention's collective-response mechanisms.
The case also raises the structural question of cloud-dependent implanted devices. The current generation of implants depend on cloud services for clinical management, firmware updates, and analytics. This dependency exposes patients to risks that are categorically different from those facing patients with self-contained legacy devices. The Convention should address the dependency directly: implanted devices should be required to remain functional in their core medical operation if cloud connectivity is lost; firmware updates should be subject to specific safeguards; the architectural choice to depend on cloud services should require regulatory approval.

CHAPTER XV-bis-15

Future-Tech Case Four — The Trans-human Inequity (2038)

[This is a fictional story based on a fictional future breakthrough.]
By 2038, a combination of consumer-grade augmentation devices and advanced inference models has produced what is generally described as the "working augmented": adults who use neurotechnology continuously and whose professional output is substantially enhanced by it. Estimates of the size of this population range from 60 to 120 million worldwide. The technology is not cheap; the working augmented are concentrated in high-income occupations in high-income countries.
A major US technology employer publishes a job posting for a senior engineering role that specifies, among the qualifications, "familiarity with continuous-augmentation workflows and demonstrated capacity for augmented collaborative work." An unsuccessful applicant, Tomas, who is not augmented and who does not wish to undergo augmentation for religious reasons, files a discrimination claim. The claim is novel: discrimination based on lack of

augmentation is not, in 2038, clearly addressed by any extant federal anti-discrimination statute, although some state statutes have begun to address it.

The legal frames available include Title VII of the Civil Rights Act (which prohibits employment discrimination on certain enumerated bases, including religion), the Americans with Disabilities Act (which prohibits discrimination against disabled persons), state-level laws (which in some states have been amended to include augmentation status as a protected category), and the emerging international jurisprudence on the right to be unaugmented.

What existing law can already do: Tomas's religious-discrimination claim under Title VII may have merit if the augmentation requirement substantially burdens his religious practice. The ADA may apply if Tomas's unaugmented state is treated as a disability by the employer, although this is a stretch. State law may provide the most direct remedy in the states that have begun to recognise augmentation status.

Where existing law falls short: the underlying structural question — whether the labour market can or should require augmentation as a condition of professional advancement — is much larger than Tomas's individual case. As augmentation becomes more pervasive, the unaugmented face cumulative disadvantage that no individual discrimination claim can address. The right to remain unaugmented, articulated in Chapter 48 and Chapter 56, is at stake; it is being eroded by market forces rather than by deliberate discrimination.

Where new jurisprudence or legislation is required: federal-level recognition of augmentation status as a protected category, with specific provisions ensuring that augmentation is not required as a condition of employment except in narrow, specifically-justified categories (e.g., particular safety-critical roles where the augmentation is itself a safety device). The proposed Convention should include such provisions at international level. National law should establish public mechanisms ensuring equitable access to augmentation for those who wish it, while preserving the right not to undergo augmentation.

The deeper question is the one this book has returned to throughout: what kind of society do we wish to construct? A society stratified between the augmented and the unaugmented, with the labour market and other structures favouring augmentation, is one trajectory. A society with deliberate equity provisions, with augmentation broadly available but never required, is another. Tomas's case is the kind of case that will arise repeatedly through the late 2030s and 2040s as the trajectory is determined by the cumulative effect of individual cases and collective policy. The choice is not yet made. The argument of this book, and the model legislation of Chapter 67-bis, is for the equity trajectory.

PART XVI

A Practical Toolkit for Citizens

The preceding parts of this guide map a terrain. This Part offers tools to navigate it. The chapters that follow are practical guidance for citizens, affected individuals, and local communities seeking to engage with the neuro-rights project at the scale of personal and civic action. The advice is concrete, accessible, and intended for sustained use rather than single reading.

CHAPTER 79

A Toolkit for the Concerned Citizen

Most readers of this book are not legislators, scientists, or activists by profession. They are concerned citizens who wish to understand what they can do, in their own lives, to defend neuro-rights and support the broader movement. This chapter offers practical guidance.

Personal Information Hygiene

The foundation of personal neuro-rights practice is information hygiene: understanding what information about yourself is being collected, by whom, and for what purposes. Several specific steps apply:

Audit consumer devices for neural-monitoring capability. EEG headbands, biosensing earbuds, smart watches with neural-signal capability, and meditation devices may collect neural data. Review the data collection and sharing policies of each. Where uncertain, assume that data is being collected and potentially shared with third parties.

Treat neural data requests with extreme caution. Where a service — medical, recreational, educational — requests neural data collection, ask: What specific data is collected? How is it stored? Who has access? How long is it retained? What inferences are drawn? Can it be sold or shared? What happens if you withdraw consent? The answers should be specific, in writing, and acceptable. If they are not, decline the service.

Apply the same scrutiny to workplace and educational deployments. Where an employer or school proposes neural monitoring, request full documentation, seek independent expert advice, and consult with worker or parent representative bodies. Do not accept neural monitoring as a condition of employment or enrolment without examining alternatives.

A Personal Neural Data Checklist

- Inventory all devices that may collect neural or biosignal data
- Review each device's privacy policy with specific attention to neural data
- Disable optional neural data collection where possible
- Request copies of neural data held about you by service providers
- Request deletion of neural data when ceasing service use
- Maintain a record of consents given, modifications, and withdrawals
- Periodically review and update consents as practices change

Civic Engagement

Personal hygiene addresses individual exposure. Civic engagement shapes the broader environment in which neurotechnology is developed and deployed. Several forms of engagement are accessible to ordinary citizens:

Contact elected representatives. Specific, written communication about neuro-rights to elected representatives is more impactful than most citizens realise. A handful of substantive constituent letters on a topic can move a legislator's priorities. Effective letters identify a specific issue, request a specific action, and explain the constituent's personal connection. They should be polite, brief, and specific. Form letters and petitions have less impact than individual messages.

Engage local government. Many neuro-rights issues touch on local government competences — schools, hospitals, police departments. Attendance at school board meetings, hospital community boards, and police oversight commissions provides opportunities to raise neuro-rights questions in concrete local contexts.

Support advocacy organisations. Several organisations work on neuro-rights, ranging from established human rights groups expanding into the field to specialised neuro-rights foundations to grassroots advocacy networks. Financial support, volunteer time, and amplification of campaigns through personal networks are all valuable.

Participate in public consultations. When governments, regulatory bodies, and standards organisations consult on neurotechnology issues, individual citizen responses are often welcomed and frequently influential. Public consultations on EU AI Act implementation, national data protection frameworks, and international human rights instruments have all included citizen input.

Professional Engagement

Many readers will have professional positions that intersect with neurotechnology directly or indirectly. Specific professional engagement opportunities exist:

Healthcare workers can advocate within their institutions for protocols that distinguish technological from psychiatric causation, support cause-neutral patient care, and integrate neuro-rights considerations into clinical practice.

Educators can incorporate neuro-rights content into curricula at appropriate levels, from primary school media literacy to university courses in law, ethics, and neuroscience. They can advocate against neural surveillance of students.

Lawyers can take pro bono neuro-rights cases, develop expertise in the emerging legal terrain, support strategic litigation, and engage in bar association activities on neuro-rights questions.

Technologists working in or adjacent to neurotechnology bear particular responsibility. They can refuse to work on systems lacking adequate human rights safeguards, advocate within their employers for ethical practices, and report serious wrongdoing through appropriate channels.

Researchers can address neuro-rights questions in their work, include human rights impact assessment in research protocols, engage with affected communities in research design, and publish in venues that reach policy audiences.

Journalists can pursue investigative reporting on neurotechnology deployments, train in the technical complexities of the field, and develop sources within companies and government agencies.

Community Organising

Beyond individual action, neighbourhood and community organisation has historically been the foundation of human rights movements. Several community-level activities support neuro-rights:

Reading groups and community discussions to develop shared understanding of the issues. The technical complexity of neurotechnology requires sustained engagement to develop confident understanding; reading groups create the conditions for this engagement.

Support networks for affected individuals. Self-identified Targeted Individuals, neurotechnology research subjects with unresolved concerns, and others affected by neurotechnology benefit substantially from peer support that is informed about the issues.

Local public events: forums, debates, film screenings, exhibitions. Cultural engagement on neurotechnology issues builds broader public awareness and creates social space for the issues to be discussed.

Coalition-building with adjacent movements: data privacy advocates, disability rights organisations, civil liberties groups, peace movements. Each has substantive overlap with neuro-rights and may welcome collaboration on shared concerns.

Maintaining Perspective

A final point of practical guidance concerns the maintenance of personal balance. Engagement with neuro-rights issues exposes the citizen to difficult material: documented abuses, complex technical threats, and the profound implications of technological developments for human autonomy. Sustained engagement requires emotional resources.

Several practices support sustainable engagement:

- Limit consumption of frightening material. Daily immersion in worst-case content produces despair without action.
- Focus on concrete, achievable actions. Specific commitments — one letter, one conversation, one meeting — are sustainable. Vague aspirations are not.
- Maintain communities of practice. Engagement with others working on the same issues provides emotional support and practical solidarity.
- Acknowledge uncertainty without paralysis. Neuro-rights involve genuine scientific and political uncertainty; effective action does not require resolution of every contested question.

- Sustain joy in everyday life. The work of defending neuro-rights serves the protection of ordinary human experience. Cultivating that experience for oneself is consistent with, not opposed to, the larger project.

The Collective Dimension

The decisive question is not what any one citizen can do, but what citizens together can do. The history of human rights demonstrates that significant change emerges from the aggregation of individual actions taken consistently over time, in coordination with others.

Each letter written, conversation held, vote cast, and donation made contributes to a collective force that, accumulated across millions of citizens, shapes the trajectory of technology and the law. The hopeful possibility is that the same collective force that has, in previous generations, secured emancipation, suffrage, civil rights, and labour protections can in this generation secure the rights of the human mind. The work begins with individual citizens choosing engagement over passivity.

CHAPTER 80

A Documentation Guide for Affected Individuals

For those who believe they may be experiencing technological targeting — whether confirmed, suspected, or contested — careful documentation is the foundation of any subsequent legal, medical, or advocacy response. This chapter offers practical guidance on what to document, how to document it, and how to preserve evidence in ways that may eventually be useful.

Purpose and Scope

This guidance is offered to assist individuals who already believe they are experiencing technological targeting and wish to document their experiences responsibly. It does not endorse any particular causal theory of their experiences. The same documentation practices serve those whose experiences ultimately have other explanations — careful records of one's own experience are valuable in any case.

Why Documentation Matters

Several distinct purposes are served by careful documentation:

For potential legal proceedings, contemporaneous records carry far more evidentiary weight than after-the-fact reconstruction. A diary entry written the day of an event, with consistent timestamps, is harder to dismiss than testimony about events recalled years later.

For medical evaluation, detailed symptom records support meaningful differential diagnosis. A clinician evaluating a patient with documented patterns of physical symptoms, with specific timing and triggers, has materially more to work with than one presented with general complaints.

For advocacy purposes, well-documented cases support broader patterns of evidence about technological targeting. Individual reports contribute to collective documentation that, in aggregate, becomes difficult to dismiss.

For one's own cognitive orientation, documentation provides a stable record against which to assess one's memory. Memory of distressing events is unreliable; contemporaneous written records provide stable ground from which to evaluate ongoing experience.

What to Document

Comprehensive documentation covers several categories:

Physical symptoms. For each symptom episode: date, time of onset, duration, location of body affected, character of the symptom (pain, heat, pressure, etc.), what the person was

doing immediately before, environmental context (location, nearby electronics, other people present), and any circumstances that appeared to modulate the symptom (movement, shielding, time of day).

Auditory phenomena. For each episode: date, time of onset, duration, perceived location of the source (inside head, external direction), content of any voices or messages, identifiable speakers or characterizations, content that could be verified externally (references to events the reporter could not know by ordinary means), and circumstances during which the phenomenon was reduced or absent.

Sleep disturbance. Daily sleep diary including bedtime, lights-out time, sleep onset time, awakenings (number, timing, apparent causes), morning waking time, total sleep estimated, sleep quality rating, and any specific phenomena (sudden awakenings, dreams of particular character, physical sensations).

Cognitive interference. Episodes of attention disruption, memory interference, intrusive thoughts, mood shifts that appear non-organic, and apparent external suggestion. Specific contexts in which these occur and any apparent triggers.

Surveillance evidence. Specific incidents that suggest surveillance: strangers approaching with apparent knowledge of recent events, unexpected service interactions, vehicles or persons appearing repeatedly in different contexts, electronic anomalies, unexplained property disturbances, and verified or suspected eavesdropping.

Coordinated harassment events. Specific incidents that suggest coordination: multiple unrelated individuals making similar comments, scripted-feeling interactions, environmental manipulation (sudden changes in store music, deliberate near-collisions), and patterns of behaviour that recur across locations.

Triggers and correlations. External events that appear to correlate with changes in targeting intensity: legal filings, public statements, contact with specific individuals, anniversaries of significant events. Documentation of these correlations supports later analysis of patterns.

How to Document

The practical mechanics of documentation matter substantially:

Use a stable system. A dedicated notebook, a digital diary app, or a structured document with consistent entry format. The choice matters less than the consistency. Multiple incompatible records are difficult to use later; a single consolidated record is far more useful.

Timestamp every entry. Date and time of writing, separately from date and time of the event being recorded. This distinction matters for legal evidentiary purposes — a contemporaneous entry made on the day of the event is more probative than a record made later.

Be specific. “Burning sensation on left arm, lateral surface, from elbow to wrist, beginning at 14:23, lasting approximately 11 minutes, while seated at desk facing south, with mobile phone 50cm away on left” is documentation. “Felt burning” is not.

Distinguish observation from interpretation. Record what you observed without entangling it with what you believe it means. “I heard a male voice saying ‘Open the letter’ while no one was in the room” is observation. “They are reading my mail” is interpretation. Both may be appropriate to record, but separately.

Preserve corroborating evidence. Photographs of physical effects (with timestamp metadata preserved), audio recordings of audible phenomena (with metadata preserved), screenshots of unusual electronic events, copies of relevant correspondence. Store these in organised folders linked to the diary entries.

Maintain off-site backup. Documents, photographs, and recordings should be backed up to off-site storage that the reporter alone controls. Multiple copies in physically separated locations provide resilience against device loss, theft, or compromise.

Consider witnesses. Where events occur in the presence of others, ask for their observation in writing. Witness statements signed and dated are valuable evidence and difficult to dismiss.

What to Avoid

Several documentation practices are counterproductive:

- **Exaggeration.** Overstated documentation undermines credibility. Conservative, precise documentation is more useful than dramatic claims that cannot be substantiated.
- **Speculation as fact.** Recording inferences as if they were observations weakens the documentation. Keep observation and inference distinct.
- **Selective recording.** Recording only confirming events while omitting disconfirming ones produces unreliable records. Document quiet periods, absent symptoms, and events that do not fit one’s theory along with those that do.
- **Public over-disclosure.** Posting detailed personal documentation to social media exposes information to potential adversaries and may reduce its evidentiary value through cross-contamination with other accounts.
- **Reliance on insecure storage.** Cloud storage controlled by major technology companies may be subject to compromise by the same actors implicated in alleged targeting. End-to-end encrypted storage, with keys controlled by the reporter, is preferable.

Medical Documentation

Establishing a relationship with a sympathetic primary care physician is valuable. Regular documentation of physical symptoms in the medical record creates contemporaneous third-party records that support both medical care and potential later proceedings. The physician need not endorse any technological theory; the physician’s role is to document what the patient presents and to investigate medical causes thoroughly.

Specific medical documentation to seek:

- Photographs of unexplained physical findings (skin lesions, burns, marks)
- Sleep studies where sleep disturbance is significant
- Audiology evaluation where auditory phenomena are reported
- Neuroimaging (EEG, MRI) where neurological symptoms are present
- Bloodwork to rule out medical causes (thyroid, autoimmune, infection)
- Documentation of any treatments tried and their effects

Legal Documentation

Reports to police and other authorities should be documented even when the response is unsatisfactory:

Make written reports rather than relying on verbal complaints. Where required to make a verbal complaint, follow up with written confirmation. Request reference numbers for all reports. Request copies of any documents created. Document the response received — even refusal to investigate is itself a documented fact relevant to subsequent advocacy and litigation.

Consider engagement with human rights bodies, even where they have no enforcement capacity. Submission to the UN Special Procedures, regional human rights commissions, or domestic ombudsmen creates official records of the complaint, regardless of immediate outcome.

The Long Game

Documentation accumulates value over time. A single dated incident is one data point; ten years of consistent documentation forms a pattern. The legal, medical, and political contexts in which the documentation may eventually be useful do not yet fully exist. The frameworks for recognising technological targeting are being built. The documentation undertaken today by affected individuals will, in many cases, become evidence usable only in those future frameworks.

This long-term horizon should inform the documentation practice. The point is not immediate vindication but sustainable, accurate, defensible records that retain their value as institutional capacity develops. Patient, methodical documentation is a form of long-term resistance against institutional dismissal — a way of insisting, in writing, that one's experience is real and demands eventual answer.

CHAPTER 81

Building Local Networks of Solidarity and Mutual Aid

Individual action and global advocacy each have their place. Between them lies a critical level often overlooked: the local network. This chapter examines the role of local networks in defending neuro-rights, supporting affected individuals, and creating the social infrastructure required for sustained movement-building.

Why Local Networks Matter

Three reasons make local networks essential to the neuro-rights project.

First, affected individuals need in-person support that online communities cannot fully provide. The isolation experienced by self-identified Targeted Individuals — family ruptures, social withdrawal, institutional dismissal — is compounded by reliance on virtual support alone. Local networks provide the embodied solidarity of regular in-person contact: a coffee with a peer who understands, a friend who can attend a medical appointment, a witness who can corroborate an event.

Second, political action requires geographic concentration. Local government, local media, and local political organising all benefit from concentrated participation by people who share concerns. A single citizen contacting one council member is one voice; twenty constituents from the same neighbourhood attending a meeting together is a force.

Third, local networks are resilient against challenges that affect virtual organising. Platform restrictions, algorithmic suppression, and surveillance of online activity all reduce the effectiveness of purely virtual organisation. In-person networks operating in physical space are less subject to these limitations.

What a Local Network Looks Like

A functional local neuro-rights network can take many forms, but typically includes several core elements:

A regular meeting cadence — weekly, biweekly, or monthly. Predictability matters more than frequency; a reliable monthly meeting is more sustainable than ambitious weekly meetings that lapse.

A small core group of committed participants — typically three to ten people — who provide continuity. Larger circles of less-committed supporters can rotate through; the core sustains the network.

A clear purpose statement that informs prospective participants what the network is and is not. A network that combines peer support for affected individuals with policy advocacy with public education can do all three if the purpose is clearly articulated; ambiguity about purpose produces friction.

Defined roles, even informal ones. Someone takes responsibility for venue arrangements, communications, and follow-through. Without role clarity, this work either does not happen or falls unsustainably on one person.

Documentation of decisions and actions. Meeting notes, action items, and a record of network activity preserve institutional memory and support new participants in understanding the network's history and direction.

Founding a Network

For those interested in founding a local network, the practical steps are accessible:

60. **Identify potential participants.** Start with one or two existing contacts who share concerns. Personal invitation is more effective than open call. The initial group establishes culture; subsequent participants are recruited into an existing pattern.
61. **Hold an exploratory meeting.** A coffee shop, a meeting room at a library, a private home. Two hours, with introductions, shared concerns, and a discussion of whether a continuing network would be valuable.
62. **If interest exists, establish basic structure.** Meeting frequency, communication channels, purpose statement. Keep it simple; structure can elaborate as needed.
63. **Hold the next meeting.** Continuation, not perfection, is the founding requirement. A network that meets twice has demonstrated more than one that meets once at high quality.
64. **Grow gradually.** Invite new participants through existing members. Avoid rapid growth that exceeds the network's organisational capacity. A small functioning network is more valuable than a large struggling one.

Safety Considerations

Networks addressing neuro-rights, particularly those involving self-identified targeted individuals, require attention to safety considerations that ordinary community groups may not face:

Information security. Meeting times, locations, and participant identities should be communicated through channels appropriate to the network's threat model. End-to-end encrypted messaging is the minimum standard. Meeting locations should be reasonably private.

Vetting of new participants. Networks attract a range of participants, including some whose presence may be disruptive or, in some cases, intentionally so. Existing members should personally vouch for new participants, and a probationary period before full participation is reasonable.

Mental health protocols. Networks supporting affected individuals will encounter participants in mental health crisis. Clear protocols for support, escalation, and referral to professional resources are essential. The network is not a substitute for professional mental health care.

Boundaries on causal claims. Networks should accommodate participants with varying beliefs about technological causation. Insistence that all participants share a specific causal theory produces division; cause-neutral validation of distress, combined with support for individual interpretations, is more durable.

Avoidance of personal risk-taking. Networks should not encourage participants to engage in confrontational or risky behaviour. Direct confrontation of alleged perpetrators, unauthorised surveillance of suspected sites, and similar tactics expose participants to legal and physical risk without commensurate benefit.

Connecting to Larger Movements

Local networks gain force through connection to larger movements:

Affiliation with national or international advocacy organisations provides resources, training, and amplification. Most major neuro-rights organisations welcome local affiliates and provide support for their development.

Participation in coordinated campaigns multiplies local impact. A letter campaign by twenty local networks across a country has substantially greater impact than the same letters from dispersed individuals.

Sharing of practices across networks accelerates organisational development. Networks face common challenges; cross-network communication accelerates the spread of effective practices.

Connection to adjacent movements expands capacity. Local networks focused on neuro-rights benefit from connection to data privacy advocates, disability rights organisations, civil liberties groups, and other allied movements operating in the same geography.

The Long Project

The construction of local networks is a long project. Networks take years to mature, decades to accumulate influence, and generations to transform political cultures. The neuro-rights movement is in its early phases; the local networks being founded today are establishing the infrastructure on which future advocacy will depend.

This long horizon is not discouraging but clarifying. The work of founding and sustaining a local network is not contingent on immediate political success. It is foundational to a sustained movement that, over the timescales of human rights development, has substantial prospect of achieving meaningful change.

The reader who founds or joins a local network today contributes to a project that may not bear fruit in their own lifetime. This is the ordinary condition of consequential civic engagement. The civil rights movement, the suffrage movement, the labour movement, the disability rights movement — each was built by people who could not see the full results of their work. The neuro-rights movement, in our generation, is building on the same foundation: patient, local, sustained, and grounded in the dignity of the people most affected.

PART XVII

Conclusion — The Choice and the Call

We have traced a trajectory that begins with laboratory curiosity and ends with the weaponization of the inner mind. The technologies described in these pages — from EEG headbands sold on Amazon to directed-energy systems developed under classified budgets — are not hypotheticals. They exist. They are deployed. And they are operating in a legal vacuum that has persisted because the pace of innovation has outstripped the capacity of democratic institutions to respond.

The core argument of this guide is simple: the last private place — the human mind — is no longer inviolable. The convergence of neurotechnology, artificial intelligence, and covert weapons programs has created a situation where thoughts can be read, emotions manipulated, and behavior influenced without the consent or knowledge of the individual. The five neuro-rights — mental privacy, personal identity, free will, fair access to cognitive enhancement, and protection from algorithmic bias — are not abstract philosophical concepts. They are the minimum conditions for human autonomy in the 21st century.

The evidence presented here is drawn from multiple sources: UN reports and resolutions, national legislation from Chile to Colorado, declassified documents tracing research from MKULTRA to DARPA's N3 program, technical analyses of directed-energy mechanisms, and the testimony of thousands of individuals worldwide who report systematic technological targeting. The consistency of these accounts — across cultures, across decades, across the boundary between classified and public knowledge — demands that we treat the phenomenon as real.

The response to date has been fragmented. Chile amended its constitution. Colorado and California protected neural data. UNESCO adopted an ethical recommendation. The UN Human Rights Council commissioned a study. These are important steps, but they are not sufficient. No binding international instrument prohibits the weaponization of neurotechnology. No oversight mechanism has the technical capacity to verify compliance. No victim recognition framework exists for those who have suffered technological abuse.

The TI One Voice Declaration, reproduced in part in these pages, represents one attempt to bridge this gap. It provides a legal framework for victims to document their experiences, connect with others, and demand investigation. It is a tool — but tools require users, and users require institutions that will take them seriously.

The choice before us is stark. In the next five years, the decisions made by legislators, judges, and citizens will determine whether neurotechnology becomes a tool for human liberation or an instrument of unprecedented control. The technology will not wait. The market will not pause. The classified programs will not be voluntarily disclosed.

What remains is the question of whether democratic societies can summon the will to protect the one place that has always been private — the space between our ears. The answer will be written not in treaties alone, but in the willingness of ordinary people to demand accountability, of scientists to speak truth to power, and of legislators to act before the next scandal makes action inevitable.

The battle for the last private place has already begun. This guide is a map. The terrain, however, must be crossed by those who read it.

CHAPTER 82

The Choice Before Us

We stand at a crossroads unprecedented in human history. For the first time, technology has developed the capacity to reach directly into the inner sanctuary of the human mind — to read its activity, to interpret its meaning, to modify its function. The UN Human Rights Council Advisory Committee, in its landmark report A/HRC/57/61 (August 2024), articulated this with stark clarity: neurotechnologies "allow a direct connection to be established between the human brain and external devices and thus provide a gateway to interfere with mental and cognitive functions."

The report identified five characteristics that make neurotechnology uniquely disruptive. It can expose cognitive processes. It can directly alter mental processes and thoughts. It can bypass conscious control or awareness. It can enable non-consensual external access to thoughts, emotions, and mental states. And it collects, analyses, and processes "large personal datasets of a highly sensitive nature" — neurodata that reveal aspects of identity and personality, including sexual orientation, personality traits, cognitive performance, intentions, beliefs, and emotions.

These are not hypothetical future capabilities. The report notes that "neuroimaging is increasingly explored and can be employed to make inferences about individuals on a large spectrum of mental states, including memories, semantic knowledge, emotions, dreams, inner speech and intentions." The convergence with artificial intelligence amplifies these capabilities, while introducing "algorithmic bias and lack of traceability and transparency."

The industry is growing at breakneck speed. Private investment in neurotechnology increased from $331 million to $7.3 billion in just ten years. Nearly half of all worldwide patent applications are registered in the United States. Major technology companies are racing to place consumer neurotechnology products on the market. And yet, as the UN report warns, "the consumer sector remains highly deregulated and is characterized by rapid expansion." The report continues: "The possibility that, in the coming years, those products with inadequate safety measures and unclear or underestimated human rights risks may be widely commercialized is real."

The human rights implications are profound. The report identifies freedom of thought as the right that "comes to the fore" when neurotechnologies have the ability to interfere with the human mind. This right protects the *forum internum* — "the inner, psychological realm of the person, where thoughts and convictions are formed." It is absolute. It does not allow for any interference whatsoever. And yet, neurotechnologies can serve as "mindreading-like" devices that make inferences about individuals, elaborate "sophisticated individualized psychological profiles," and use them to predict behaviour or modify thoughts.

The right to privacy faces new threats. Neurodata provide "highly sensitive information" that can be sold to third parties without authorization. Consumer companies are already "collecting unknown quantities of their users' personal data, which can be stored or even sold with the obvious expectation that they will become decipherable in the future." The right to

personal integrity is engaged when neurointerventions "produce effects on the body and the psychological sphere" — when unauthorized access to mental activity aims to "intrude into mental processes, alter mental functioning or seek to manipulate."

And beyond the framework of established human rights law lies a darker reality. The UN report acknowledges that neurotechnologies "can be used to interfere and manipulate individuals" through neuromodulation devices that alter mental processes "in ways similar to 'brainwashing.'" The report references "neuromarketing" strategies that "condition the forming of opinions" and influence decision-making. It notes that "neurogaming" devices can "drive compulsive or addictive use, particularly when neural reward circuits target altering consumption behaviours."

But the weaponization of neurotechnology extends far beyond commercial manipulation. The testimony of thousands of individuals worldwide — documented in the TI One Voice Declaration and supported by historical records of classified programs from MKULTRA to DARPA's N3 program — describes systematic technological targeting: remote neural surveillance, synthetic telepathy, cognitive and emotional manipulation via electromagnetic means. These accounts, consistent across cultures and decades, describe a form of human rights violation that existing legal frameworks are ill-equipped to address.

As one submission to the OHCHR from a Belgian researcher and targeted individual states: "Current transitional justice mechanisms are ill-equipped to address crimes that leave no physical evidence visible to conventional forensics, are perpetrated remotely across jurisdictions, employ technologies officially denied or classified, and create isolation through systematic gaslighting by institutions."

The UN report calls for action. It states: "Although human rights violations may happen as a consequence of actions and misuses, they are also the result of lack of preparedness. Where areas not sufficiently covered by existing standards do emerge, action is needed to prevent violations, misuses and abuses." It recommends developing "context-specific standards" to better define the scope of rights being engaged, and reinforcing "the institutional structure to ensure that prevention and mitigation policies are effectively implemented."

The choice before us is not between action and inaction. Inaction is itself a choice — one that cedes the terrain to market forces and classified programs operating without democratic oversight. The UN report warns that "without the necessary guardrails, the industry will continue growing unfettered in the same direction: prioritizing profitability and convenience over ethical and human rights considerations."

The next five years will be decisive. The technology will not wait. The market will not pause. The classified programs will not be voluntarily disclosed. What remains is the question of whether democratic societies can summon the will to protect the *forum internum* — the last private place — before its walls are breached beyond repair.

CHAPTER 83

A Call to Action — What You Can Do

This book closes where it began — with the human mind and the threats now arrayed against it. The preceding eighty-two chapters have laid out, in detail, the science, the law, the history, the politics, and the lived experience. What remains is the question every reader must ask: *and now*, what will I do?

The answer is not the same for everyone. A neuroscientist will act differently from a city councillor, a teenager differently from a retiree, a Brussels resident differently from a citizen of São Paulo or Manila. But the structure of action is the same. There are tiers of engagement that cumulatively form a single defence; there is a small set of concrete moves any reader can begin this week; and there is a shared set of principles that allow individual actions to compound into collective change. This chapter supplies all three: a framework, a checklist, and a closing argument for why this matters.

A Pyramid of Engagement

Engagement is not a single decision but a spectrum. At the broadest base sits awareness: knowing what neurotechnology is, how it works, and what is at stake. Above that sits self-protection: auditing the devices in one's own life and demanding that the data they collect remain under one's control. Above that, public voice: speaking to legislators, journalists, and the institutions one already belongs to. Above that, organising: building or joining the local groups through which sustained advocacy is possible. At the apex, leadership: standing for office, litigating, founding civil-society organisations, or drafting the treaties of the next decade.

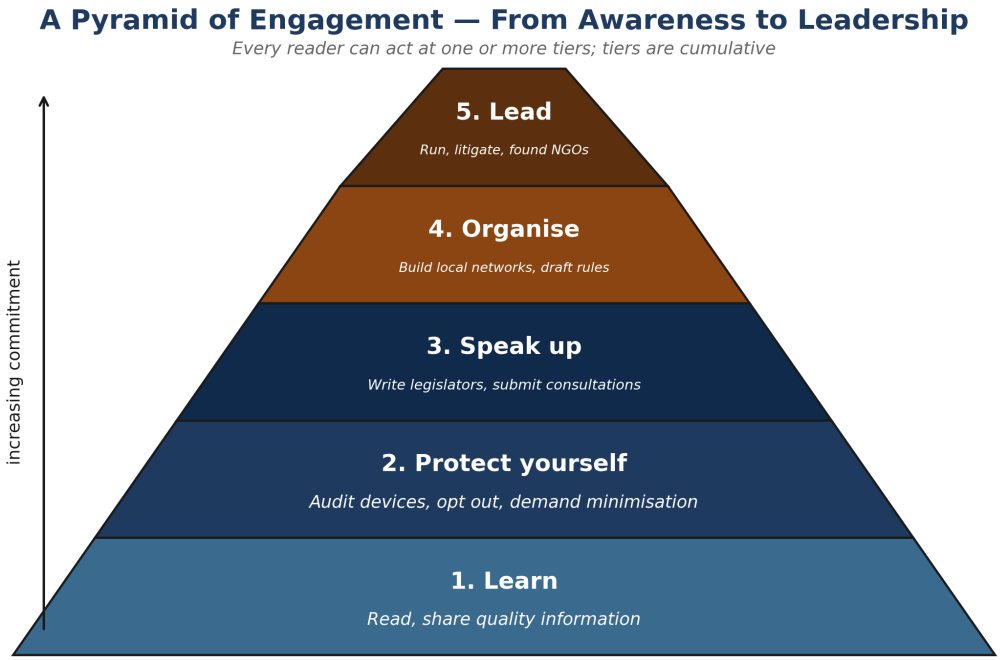


Figure 83.1. *A pyramid of engagement — from awareness at the base to leadership at the apex. Tiers are cumulative, not exclusive: a citizen who acts at one tier is more, not less, likely to also act at the tiers below.*

The pyramid is cumulative, not hierarchical. A reader who organises a neighbourhood circle is almost certainly also auditing their devices and writing to their legislators. The point of the diagram is not to rank people; it is to make visible the fact that effective movements depend on participants at every tier. A movement composed entirely of leaders will exhaust itself; a movement composed entirely of readers will achieve nothing. Mental sovereignty needs both.

Ten Minutes, Ten Hours

The question this chapter most often elicits is the practical one: *where do I start?* The answer depends on how much time the reader has this week. The decision tree below maps the choice. Whichever branch is chosen, the closing step is the same: document what you did, tell one other person, and repeat next week. That, and not any grand campaign, is what built Chile's constitutional amendment, the UNESCO recommendation, and every other concrete win this book has chronicled.

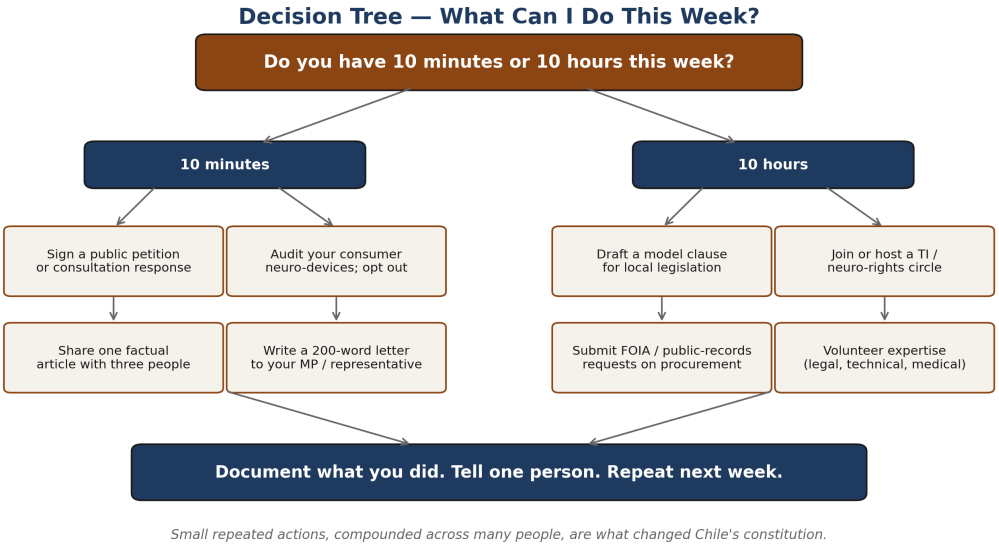


Figure 83.2. *A practical decision tree for the reader who has just finished this book. Each leaf is an action that can be completed without specialist training, by an ordinary citizen, in a single week.*

Ten-minute actions

- **Sign a petition or submit a consultation response.** Many national and international bodies have ongoing consultations on neurotechnology, AI, and data protection. UNESCO, the OECD, the European Commission, the UN Human Rights Council, and most national parliaments take written submissions from private citizens. A short, well-argued letter counts.
- **Audit your consumer neuro-devices.** If you own an EEG headset, a sleep-tracking band that monitors brain activity, a meditation app that reads physiological signals, or a VR headset with eye-tracking, check the privacy settings, read the data-sharing terms, and turn off whatever is not essential. Where opt-out is available, take it.
- **Share one factual article with three people.** Misinformation spreads not because it is powerful but because it travels. So does good information, if its carriers choose to move it. Pick one article from a credible source — not this book, not a single advocacy group, but a newspaper, journal, or institutional report — and send it to three people in your network with a single sentence of context.
- **Write a 200-word letter to your representative.** National legislators count letters. Two hundred words is enough to identify yourself, name a specific concern about neurotechnology or neural data, and ask a specific question. The reply, if any, is a public record you can share.

Ten-hour actions

- **Draft a model clause.** Many municipalities, universities, hospitals, and trade unions have policies on data protection but no specific provisions on neural data. A short, clearly drafted clause — defining neural data, requiring opt-in consent,

prohibiting downstream sale — can often be adopted by an institution that would not have written its own.

- **Join or host a local neuro-rights circle.** The model is simple: a recurring meeting of five to fifteen people who share the concern, exchange information, and coordinate small public actions. Local circles are the connective tissue of every successful rights movement of the twentieth century, and they remain the most underused tool in the neuro-rights toolkit.
- **Submit a freedom-of-information request.** Public procurement records, research-grant disclosures, police equipment inventories, and university ethics-board approvals are subject to FOI laws in most democracies. A targeted request — for example, asking what neurotechnology products a given police department has procured — produces information that no journalist will otherwise gather.
- **Volunteer expertise.** If you are a lawyer, an engineer, a clinician, a translator, a designer, or an accountant, you possess a skill that one of the small civil-society organisations working on these issues can use. Send them an email offering a specific number of hours per month. The answer will not always be yes, but it will sometimes be yes.

Protecting Yourself First

Effective advocacy is impossible from a posture of self-neglect. The reader who is exhausted, frightened, or under active surveillance cannot draft model legislation. Self-protection is therefore not a retreat from advocacy; it is its precondition. The following list is a starting point. It is not exhaustive, and the appropriate measures depend on individual circumstances and threat models.

Personal Mental-Hygiene Checklist

• Audit the consumer devices that touch your body or capture your physiological state. Disable non-essential data collection. • Use end-to-end encrypted messaging for sensitive communication. • Keep at least one offline notebook for reflective writing that no platform can read. • Maintain ordinary sleep, exercise, and social contact — the foundation of cognitive resilience. • Document any pattern of experience that troubles you, with dates, times, and concrete details, before reaching conclusions. • If you believe you are being targeted, consult both a trusted clinician and a trusted lawyer; the two perspectives are complementary, not contradictory. • Build a small circle of people who know your situation and whose judgement you trust.

No One Acts Alone — The Ecosystem of Defence

Mental sovereignty cannot be secured by any single actor. Citizens cannot pass laws; legislators cannot conduct independent science; journalists cannot litigate; industry cannot self-regulate without external oversight. The diagram below names the six categories of actor on whom the defence of the inner realm depends. The reader of this book occupies the “citizen” role by default and may occupy others by profession, training, or election.

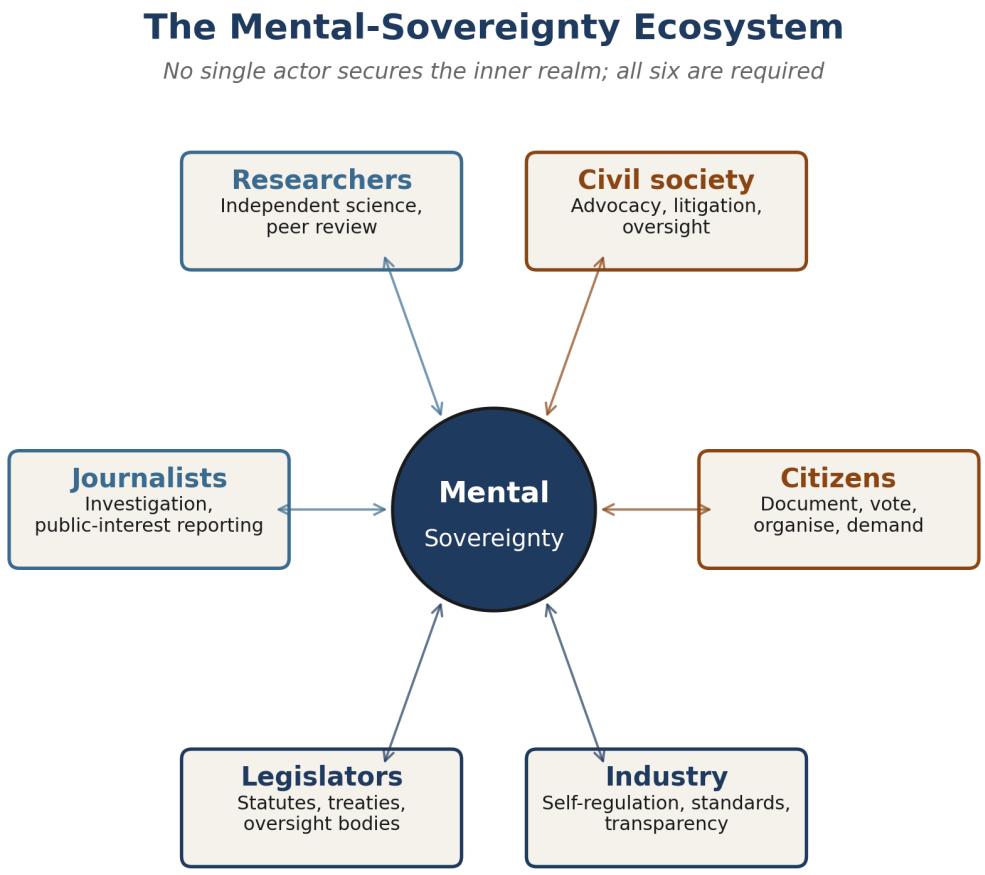


Figure 83.3. *The mental-sovereignty ecosystem. No single actor can secure the inner realm; the work depends on the interaction of all six categories. Each reader belongs to at least one.*

The ecosystem view also clarifies why some forms of action are higher-leverage than others. Persuading a journalist to investigate a particular procurement is, in expectation, more consequential than persuading a single citizen of the same fact — because the journalist's output reaches thousands of citizens. Persuading a legislator to introduce a bill is more consequential still, because a bill survives the legislator's term and continues to do work. The reader who finds themself with access to a journalist, a legislator, a researcher, or a regulator should consider that access a public trust, and use it accordingly.

Five Principles That Sustain Engagement

Across the diverse activists, researchers, and citizens whose work this book draws on, five working principles recur. None is original to this volume; all are worth restating because each, in our experience, marks the difference between sustained engagement and burnout.

1. Document before you advocate.

Advocacy that is not anchored in documented fact is brittle. The first move in every successful campaign of the past decade has been the assembly of a documentary record — declassified files, procurement records, peer-reviewed papers, testimony archives. Documentation is slow work, but it is the work that gives every subsequent action a foundation that cannot be argued away.

2. Local before global.

International treaties matter, but they are negotiated by national governments, and national governments respond to domestic constituencies. The reader who wants to influence the UN is well-advised to begin with their city council. Mental sovereignty, like most rights, is won locally and then exported upward through the diplomatic stack.

3. Boring before dramatic.

The actions that produced Chile's constitutional amendment were not dramatic. They were committee drafting sessions, expert testimony, line edits to legislative text, and patient correspondence with ministers. The dramatic moments — the floor vote, the signing ceremony — were the visible tip of years of ordinary, unglamorous, boring work. The reader who is willing to do boring work has a decisive comparative advantage.

4. Coalitions, not silos.

Neuro-rights overlaps with privacy advocacy, AI ethics, disability rights, patient rights, labour rights, free-speech advocacy, religious-liberty advocacy, and several traditions of academic ethics. Each of these constituencies has organisations, lawyers, lobbyists, and constituencies of its own. Allies multiply where they are sought. Coalitions of convenience — on a specific bill, a specific case, a specific consultation — are nearly always more effective than purist solo efforts.

5. Sustainability before intensity.

The reader who works fourteen hours a day for two months and then collapses contributes less than the reader who works two hours a week for twenty years. Mental-sovereignty work is generational. Pacing is not a luxury; it is a strategic requirement. Burnt-out advocates do not write the laws of the next decade.

The Choice and the Call

The choice this book has tried to make visible is now in the reader's hands. The diagram below compresses it into a single image. From the present moment, two paths fork — one toward drift, the other toward defence. Neither outcome is automatic; both require sustained collective action by millions of people whose names will never appear in a book.

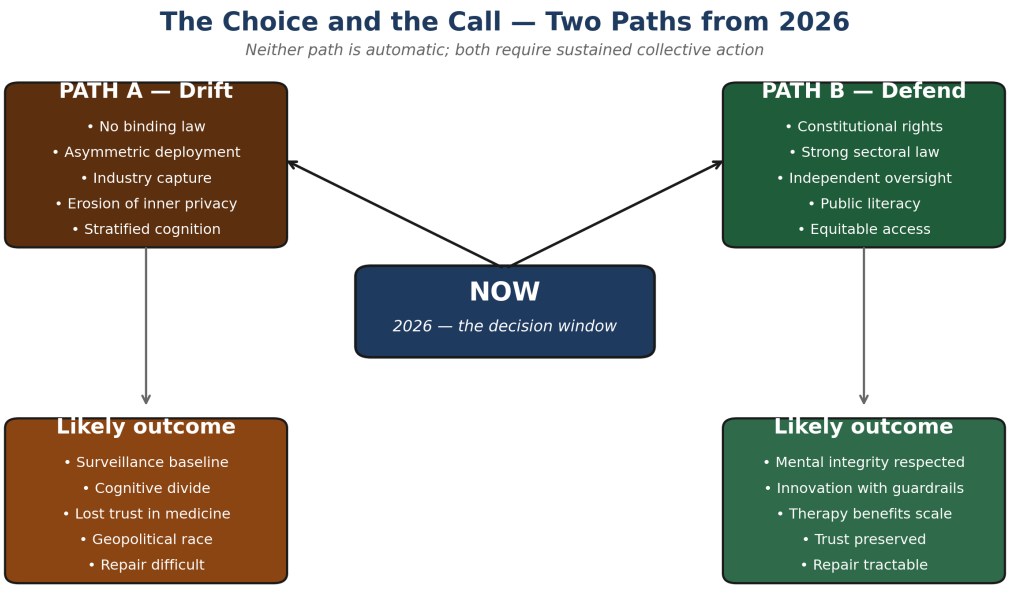


Figure 83.4. *The choice and the call — two paths from 2026. Path A (drift) requires only inaction. Path B (defence) requires a sustained collective effort across the ecosystem described above.*

Path A — drift — is the path of least resistance. It requires no decision and no action. The default rate of technological deployment, the default rate of regulatory capture, and the default rate of human attention will combine to produce it in the absence of countervailing effort. Its likely outcomes — a surveillance baseline, a cognitive divide, an erosion of medical trust, a geopolitical race — are not predictions of malice. They are the predictable consequences of the absence of deliberate choice.

Path B — defence — is the path that requires choice. Its preconditions are the topics of this book: constitutional recognition, sectoral statute, independent oversight, public literacy, equitable access, international architecture. Each precondition requires legislators willing to draft, lawyers willing to litigate, journalists willing to investigate, researchers willing to speak, industry willing to submit to audit, and citizens willing to insist. Path B is not a guarantee; it is a possibility that opens only if enough actors choose it.

The reader of this book is one of those actors. The book has tried to be honest about the difficulty: many of the actors needed are not currently in the field, many of the laws needed are not currently on the books, and many of the institutions needed do not currently exist. None of this is unprecedented. Slavery was abolished, women won the franchise, decolonisation was completed, and the bare outlines of international human-rights law were assembled — each of these within living memory, each against longer odds than the ones currently facing the neuro-rights movement. The people who made those changes were not, on the whole, more talented or better resourced than the readers of this book. They simply decided to begin, and then they did not stop.

A Final Word

The thesis of this book has been that the human mind is the last private place, and that its defence is the central political task of our generation. The argument is not a counsel of

despair. The technologies are real but young; the laws are insufficient but improving; the public is uninformed but interested. The window of decision is open. It will not remain open forever, and it will not close with an announcement. It will close one procurement at a time, one regulatory exemption at a time, one missed consultation at a time — quietly, in the absence of opposition.

The call of this book is not heroic. It is incremental. Read. Audit. Speak. Organise. Document. Repeat. The change required is not a single act of valour by a small number of remarkable people. It is a sustained act of ordinary citizenship by a very large number of ordinary people. That is the form most lasting political change has taken throughout modern history, and it is the form mental sovereignty will take if it is to take any form at all.

The work begins now. It begins with the reader of this sentence, in whatever city or village this book has reached, with whatever ten minutes or ten hours can be set aside this week. The argument is over; the action remains. The mind — your mind, every mind — deserves no less.

For everyone who has ever felt their thoughts were not their own —

and for everyone who refused to accept that they were.

PART A

Glossary of Key Terms

This glossary defines key technical, legal, and ethical terms used in this guide. The definitions are intended for general readers and prioritise accessibility over technical precision; readers seeking authoritative definitions should consult the relevant specialised literature.

Term	Definition
Algorithmic Inference	A conclusion drawn by an AI system from raw data, such as inferring emotional state, attention level, or intention from neural signals.
Brain-Computer Interface (BCI)	A direct communication pathway between the brain and an external device, either invasive (surgically implanted) or non-invasive (worn on the skin).
Cognitive Liberty	The right to self-determination of one's own mental processes, including the right to refuse neural surveillance and to use or refuse cognitive enhancement.
Cognitive Warfare	The use of technologies and techniques to alter human cognition without the target's knowledge or consent, typically in a military or geopolitical context.
Convention on the Rights of Persons with Disabilities (CRPD)	United Nations treaty (2006) addressing the rights of persons with disabilities, relevant to neuro-rights given the role of neurotechnology in disability and the autonomy concerns it raises.
DARPA N3	Next-Generation Non-Surgical Neurotechnology programme of the US Defense Advanced Research Projects Agency, developing non-invasive military brain-computer interfaces.
Deep Brain Stimulation (DBS)	A medical procedure in which electrodes implanted in the brain deliver electrical impulses to treat conditions such as Parkinson's disease, epilepsy, and severe depression.
Directed-Energy Weapon	A device that delivers targeted energy (electromagnetic, acoustic, optical) at a distance, with potential effects on biological tissue and neurological function.
EEG (Electroencephalography)	The non-invasive measurement of electrical activity of the brain through electrodes placed on the scalp.
fMRI (Functional Magnetic Resonance Imaging)	A neuroimaging technique that measures brain activity by detecting changes in blood flow.
Forum Internum	In international human rights law, the inner mental and conscience realm where thoughts and convictions are formed, traditionally considered inviolable.
Foundation Model	A large-scale AI model trained on broad data and adaptable to many downstream tasks, increasingly applied to neural data decoding.
Havana Syndrome	A set of unexplained neurological symptoms reported by US diplomats and intelligence officers, hypothesised to result from directed-energy attack.
Mental Privacy	The right to keep one's mental processes, thoughts, and neural data free from non-consensual access or surveillance.
Microwave Auditory Effect	A phenomenon in which pulsed microwaves can be perceived

Term	Definition
	as sound, providing a physical basis for some claims of “voice-to-skull” technology.
MKULTRA	CIA programme (1953-1973) of covert experimentation on human subjects, including drug administration, sensory deprivation, and behavioural modification.
Neural Data	Information collected from the brain or nervous system, including both raw signals and algorithmic inferences derived from them.
Neurodiscrimination	Adverse treatment of individuals based on their neural data, neural inferences, or neurological characteristics.
Neuro-rights	A category of human rights addressing the unique challenges posed by neurotechnology, typically including mental privacy, personal identity, free will, fair access to cognitive enhancement, and protection from algorithmic bias.
Neurotechnology	Any device, system, or method capable of measuring, recording, interpreting, stimulating, or modulating neural activity.
Neuroweapon	A device or system designed to cause harm through direct effect on the nervous system.
Personal Identity	The right to maintain the continuity of one’s sense of self in the face of neurotechnological intervention.
Targeted Individual (TI)	A person who reports being subject to non-consensual technological targeting, typically including directed-energy harassment, synthetic auditory phenomena, and pervasive surveillance.
UNESCO Recommendation on the Ethics of Neurotechnology	Global ethical framework adopted by UNESCO member states in November 2025, providing the first international standard for neurotechnology ethics.
Zersetzung	East German Stasi programme of coordinated psychological harassment of political dissidents, considered a historical analogue to some Targeted Individual reports.

PART B

Timeline of Neuro-rights Milestones

This appendix provides a chronological reference of key events, publications, and legal developments relevant to the neuro-rights field. It is not exhaustive but identifies the events most frequently cited in this guide.

Year	Milestone
1953	CIA establishes Project MKULTRA, the foundational programme of covert behavioural experimentation in the United States.
1974	US Department of Health, Education, and Welfare publishes *Individual Rights and the Federal Role in Behavior Modification*, raising early concerns about direct brain stimulation research.
1976	US Senate Church Committee hearings publicly document MKULTRA, establishing the historical baseline for institutional concern about covert neurotechnology.
1979	Soviet Union initiates major investment in psychotronics research, the foundation of subsequent Russian cognitive warfare programmes.
1980s	Frey, Lin, and other researchers document the microwave auditory effect, establishing the physical basis for later “voice-to-skull” claims.
1990s	Initial deep brain stimulation devices receive regulatory approval for medical use.
2002	US Department of Defense formalises the NBIC (Nanotechnology, Biotechnology, Information Technology, Cognitive Science) Convergence Project.
2004	US Air Force Research Laboratory publishes report *Bioeffects of Selected Non-Lethal Weapons*, documenting directed-energy effects research.
2013	US BRAIN Initiative announced; EU Human Brain Project begins.
2014	China launches major brain research initiative; private neurotechnology investment estimated at $331 million annually.
2016	Reports of Havana Syndrome among US diplomats in Cuba begin.
2017	Neuralink founded by Elon Musk; Synchron and other major BCI startups raise large funding rounds.
2018	Nita Farahany publishes *The Battle for Your Brain*; Roberto Andorno

Year	Milestone
	articulates the five core neuro-rights.
2019	Meta acquires CTRL-Labs for an estimated $1 billion, signalling major-tech entry into neural interfaces.
2021	Chile becomes the first country to amend its constitution to protect neuro-rights, adding cerebral protection to Article 19.
2022	UN Human Rights Council Resolution 51/3 (6 October 2022) requests comprehensive study of neurotechnology and human rights.
2023	Meta demonstrates intelligible speech reconstruction from MEG; private neurotechnology investment reaches $7.3 billion annually.
2024	EU AI Act adopted; UN Human Rights Council Advisory Committee publishes A/HRC/57/61 (8 August 2024); Council of Europe Framework Convention on AI adopted; Colorado and California amend privacy laws to protect neural data.
2025	UNESCO Recommendation on the Ethics of Neurotechnology adopted in November 2025 by 190+ member states.
2026	TI One Voice global declaration reaches signatories from 36 countries; expanded petitions submitted to the European Parliament and US Congress.

PART C

Country-by-Country Legislative Status

This appendix summarises the state of neuro-rights legislation in selected jurisdictions as of early 2026. The picture is rapidly evolving; readers should consult current sources for definitive information.

Jurisdiction	Status
Chile	Constitutional amendment (2021) protecting cerebral activity. Implementing legislation under development. First country with explicit constitutional neuro-rights.
Brazil	Constitutional reform proposals under discussion since 2024. Active civil society engagement.
Mexico	Neuro-rights legislation proposed in 2024. Senate hearings held; passage not yet achieved.
Spain	Charter of Digital Rights (2021) includes neuro-rights principles. EU regulatory framework applies.
European Union	EU AI Act (2024) classifies certain neurotechnology applications as high-risk. Council of Europe Framework Convention on AI (2024) provides additional principles. GDPR applies to neural data as biometric data.
United States (federal)	No federal neuro-rights legislation. FDA regulates medical neurotechnology. Defence research programmes (DARPA N3) continue.
United States (Colorado)	Amended privacy law (2024) classifies neural data as sensitive personal information requiring explicit consent.
United States (California)	CCPA amendment (2024) protects neural data and algorithmic inferences derived from it. Strongest US state-level protection.
United Kingdom	No specific neuro-rights legislation. ICO has issued guidance on biometric data. House of Lords committee report on neurotechnology (2024).
Canada	Privacy legislation (PIPEDA) applies. Specific neuro-rights legislation under discussion.
Australia	No specific neuro-rights legislation. Privacy Act 1988 applies.

Jurisdiction	Status
Japan	Brain/MINDS programme ongoing. No specific neuro-rights legislation; existing privacy frameworks apply.
South Korea	Active brain research programme. No specific neuro-rights legislation.
China	Major brain research investment. No neuro-rights protective legislation; significant state interest in cognitive monitoring applications.
Russia	Historical psychotronics research; current programmes largely opaque. No protective legislation.
India	No specific neuro-rights legislation. Digital Personal Data Protection Act (2023) provides general framework.
African Union	African Charter on Human and Peoples' Rights provides general framework. Specific neuro-rights initiatives at early stages.
Arab League	No regional neuro-rights framework. Individual member state approaches vary.
International Bodies	UN Human Rights Council A/HRC/57/61 (2024) provides analytical foundation. UNESCO Recommendation (2025) provides ethical framework. No binding international treaty exists.

PART D

Resources and Further Reading

This appendix lists key resources for readers seeking to engage further with the issues raised in this guide.

International Bodies and Reports

- UN Human Rights Council Advisory Committee Report A/HRC/57/61 on the impact of neurotechnology (8 August 2024) — the authoritative international examination of neurotechnology and human rights.
- UNESCO Recommendation on the Ethics of Neurotechnology (November 2025) — first global ethical framework.
- UN Special Rapporteur on Torture report on cyber-torture (2020) — application of torture framework to technologically mediated harm.
- Council of Europe Framework Convention on Artificial Intelligence (2024) — principles-based framework with implications for neurotechnology.
- OECD Recommendation on Responsible Innovation in Neurotechnology (2019) — early framework for ethical neurotechnology development.

Civil Society Organisations

- TI One Voice (one.witysk.org) — global initiative documenting Targeted Individual experiences and advocating legislative reform.
- The Neurorights Foundation — international advocacy for neuro-rights including the five core rights framework.
- Access Now — digital rights organisation with growing neurotechnology focus.
- European Digital Rights (EDRi) — European coalition addressing neurotechnology among other digital rights issues.
- Future of Life Institute — includes neurotechnology in broader technology ethics work.

Academic Resources

- Marcello Ienca and Roberto Andorno, “Towards new human rights in the age of neuroscience and neurotechnology” (2017) — foundational article articulating five core neuro-rights.
- Nita Farahany, *The Battle for Your Brain* (2023) — accessible introduction to cognitive liberty and neuro-rights.

- Rafael Yuste et al., “Four ethical priorities for neurotechnologies and AI” (Nature, 2017) — widely cited ethics framework.
- Stephen Rainey and Yasemin J. Erden, *Neurotechnology and the Law* (forthcoming) — comprehensive legal treatment.

For Affected Individuals

- TI One Voice support resources at one.witysk.org — includes the global Declaration, documentation guides, and connection to support networks.
- National Crisis Lines — readers in mental health crisis should contact local crisis services. In the EU, the European Emergency Number 112 connects to emergency services. The Samaritans (UK and Ireland) operate at 116 123. In the US, 988 is the Suicide and Crisis Lifeline. Other countries maintain equivalent services.
- Cause-neutral peer support networks — several organisations provide peer support without requiring commitment to specific causal theories.

Legal Resources

- EU AI Act full text and implementation guidance (available at digital-strategy.ec.europa.eu).
- GDPR full text and EDPB guidance on neural data (available at edpb.europa.eu).
- Chilean Constitution Article 19 amendment text (available through the official gazette of Chile).
- California Consumer Privacy Act amendments (2024) on neural data (available at oag.ca.gov).

Staying Current

The field is evolving rapidly. Readers seeking to stay current should consider following: academic journals including *Neuroethics*, *Nature Neuroscience*, and *AJOB Neuroscience*; civil society newsletters from the organisations listed above; and parliamentary hearings and regulatory consultations in jurisdictions of interest. Annual updates to this guide are anticipated; readers are invited to consult the publisher’s website for the most current edition.

PART E
References — Complete Bibliography
This appendix consolidates the full reference apparatus of the book, comprising: (i) the original Document Library and Knowledge Base references established for the first and second editions; (ii) the references underlying the source essay The Last Private Place; and (iii) the references for the new material added in the fourth edition. The list is intentionally comprehensive: readers preparing legislative proposals, academic articles or advocacy submissions should be able to locate the underlying source for any contested claim in the book.

Section 1 — Document Library (Original)

The following items were curated for the Document Library of the TI One Voice initiative and constitute primary source material referenced throughout this book.
Electronic Mind Control Book (TI Book, 2026-03-06) — Document Library ID 487
Legislative change petition EU (Mail-template, 2026-05-26) — Document Library ID 4141
Neurorights doc EU (Mail-template, 2026-05-27) — Document Library ID 4168
The Last Private Place (article, 2026-05-27) — Document Library ID 4162
TI One Voice Declaration - Stephane van der Aa (declaration, 2026-04-15) — Document Library ID 2565
Transhuman Reader Book (TI Book, 2026-03-06) — Document Library ID 490
Understanding the Targeted Individual Phenomenon: An Educational Brief (Personal stories, 2026-05-27) — Document Library ID 4205
United Nations states the dangers of Neurotechnologies committing Human Rights Abuses (Evidence, 2026-03-06) — Document Library ID 510

Section 2 — Knowledge Base (Original)

The following items are research papers, evidentiary documents and historical materials curated for the Knowledge Base of the TI One Voice initiative, organised by category. Files marked with a Knowledge Base ID are available through the platform at one.witysk.org.
$MP-HFM-334-KN3.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 169
119049866-Russian-and-US-Mind-Control-Technology-Psychotronic-War.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 81
14_un_special_rapporteur_submission.docx (Letter Templates, 2026-04-19) — Knowledge Base ID 450
159405705-Darpa-on-Your-Mind.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 258
1Voice Declaration - legal doc and guide v3.1b.pdf (Legal Resources, 2026-05-26) — Knowledge Base ID 482
1Voice Declaration legal doc and guide v1.1 (Legal Resources, 2026-01-28) — Knowledge Base ID 21
2EmondiPresentationPDFversion2.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 45
Anna Fubini Mind Weapon (Technology Information, 2026-01-28) — Knowledge Base ID 19
Cognitive_liberty._A_first_step_towards.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 92
CPDE_Complete_Unified_Report_20260416.pdf (TI Phenomenon Research, 2026-04-16) — Knowledge Base ID 435
ElectromagneticTorture.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 156
flyer-brussels-fr-print-finish.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 148
G2250614 (1).pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 160
Global Brain Chip and Mesogens.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 152
Havana-Syndrome-report.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 150
ico-tech-futures-neurotechnology-0-1.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 151
Legislative reform proposals (content-page, 2026-04-05) — Knowledge Base ID 426

Microwave Radio Frequencies 5G 6G Graphene Nanomaterials Technologies used in neurological warfare DeruelleF (2024).pdf (Scientific Papers, 2026-03-26) — Knowledge Base ID 388
Mind Control: From Nazis to DARPA - David Salinas Flores.PDF (Scientific Papers, 2026-03-26) — Knowledge Base ID 393
Monitoring_and_manipulating_the_human_br.pdf (Scientific Papers, 2026-03-26) — Knowledge Base ID 382
Neurowar Is Here! - Joshua D Gramm & Brian A Branagan (2021).pdf (Scientific Papers, 2026-03-26) — Knowledge Base ID 384
RAND_RR2996.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 52
The New Era of Neural Modulation Led by Smart Nanomaterials (2024).pdf (Scientific Papers, 2026-03-26) — Knowledge Base ID 389
The_Silent_Massacre_Research_Report.pdf (TI Phenomenon Research, 2026-04-02) — Knowledge Base ID 406
TI_One_Voice_DA_Study_Addendum_Milestone_Zero_Official_Recognition.docx (TI Phenomenon Research, 2026-04-04) — Knowledge Base ID 412
whitepaper_neurogress.pdf (TI Phenomenon Research, 2026-03-08) — Knowledge Base ID 42
Wireless agents for brain recording and stimulation modalities (2023).pdf (Scientific Papers, 2026-03-26) — Knowledge Base ID 390

Section 3 — Source Essay References

The book draws extensively on the source essay "The Last Private Place" by Stephane van der Aa (TI One Voice, 2026). The numbered references below are those of the essay itself and are reproduced here for the convenience of the reader.

[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).

Section 4 — References for Fourth-Edition New Material

The following references support the chapters added in the fourth edition of this book. References are grouped by the chapter in which they are most directly relevant; many are relevant to multiple chapters and are listed only once for compactness.

Chapter 18-bis — Existing Laws at Every Level

International Covenant on Civil and Political Rights, Articles 7, 17, 18, 19 (UN, 1966).
Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment (UN, 1984).
Convention on the Rights of the Child (UN, 1989); Convention on the Rights of Persons with Disabilities (UN, 2006).
Biological Weapons Convention (1972); Chemical Weapons Convention (1993); Convention on Certain Conventional Weapons (1980).
European Charter of Fundamental Rights, Articles 1, 3, 7, 8, 10 (EU, 2000).
European Convention on Human Rights, Articles 8, 9, 14; Council of Europe Convention 108+; Oviedo Convention on Biomedicine.
Pretty v. United Kingdom, ECHR (2002); V.C. v. Slovakia, ECHR (2011); Selmouni v. France, ECHR (1999); Puttaswamy v. Union of India, Supreme Court of India (2017).
African Charter on Human and Peoples' Rights (1981); Malabo Convention on Cyber-Security and Personal Data Protection (AU, 2014, in force 2023).
Arab Charter on Human Rights (2004); ASEAN Framework on Personal Data Protection (2016, rev. 2024).
Chile Law 21.383 (2021); Brazil Constitutional Amendment on Mental Integrity (2024); Spain Charter of Digital Rights (2021).
EU General Data Protection Regulation (Regulation 2016/679); EU AI Act (Regulation 2024/1689); EU Medical Device Regulation 2017/745.
US 1st, 4th, 5th, 14th Amendments to the United States Constitution; German Basic Law Articles 1, 2; French Declaration of the Rights of Man (1789).

Colorado Privacy Act amendment on neural data (2024); California CCPA amendment (2024); Montana Genetic Information Privacy Act amendment (2025); Connecticut Data Privacy Act amendment (2025).
Genetic Information Nondiscrimination Act (GINA), 42 U.S.C. §2000ff (US, 2008).
WTO Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS); General Agreement on Trade in Services (GATS); Comprehensive and Progressive Agreement for Trans-Pacific Partnership (CPTPP).

Chapter 33-bis — Where Mind-Control Risk Comes From

Neurorights Foundation. "Safeguarding Brain Data: Assessing the Privacy Practices of Consumer Neurotechnology Companies" (Genser, Damianos & Yuste, 2024).
Zuboff, S. The Age of Surveillance Capitalism. PublicAffairs (2019) — context for incentive-misalignment analysis.
BrainCo Inc. / FocusEDU classroom deployment coverage; SCMP and Wall Street Journal investigative reporting (2019-2024).
NATO Allied Command Transformation. Cognitive Warfare exploratory studies; Innovation Hub publications.
U.S. Office of the Director of National Intelligence. Assessment on Anomalous Health Incidents (March 2023).
Wittes, B. & Friedland, S. "Cognitive Warfare and Civil Liberties." Lawfare and Brookings (2023-2024).
Kreitmair, K. "Dimensions of Ethical Direct-to-Consumer Neurotechnologies." AJOB Neuroscience (2019).
Ienca, M. "Common Human Rights Challenges Raised by Different Applications of Neurotechnologies in the Biomedical Field." Council of Europe (2021).

Chapter 62-bis — The Device Makers Atlas

Medtronic Activa/Percept clinical literature; Boston Scientific Vercise programming literature; NeuroPace RNS System publications.
Pandarinath, C. et al. "High performance communication by people with paralysis using an intracortical brain-computer interface." eLife (2017) — Blackrock-relevant.
Oxley, T. J. et al. "Motor neuroprosthesis implanted with neurointerventional surgery improves capacity for activities of daily living tasks in severe paralysis." Journal of NeuroInterventional Surgery (2021) — Synchron Stentrode.
Neuralink Corporation. PRIME study public materials (2024-2026).
Precision Neuroscience. Layer 7 Cortical Interface public clinical materials.
Paradromics, Inc. Connexus Direct Data Interface public materials.
Inbrain Neuroelectronics. Graphene-based neural interface clinical materials.
Onward Medical N.V. ARC-EX and ARC-IM public clinical literature.
g.tec medical engineering GmbH product literature; InteraXon (Muse) public terms of service and developer documents; Emotiv public materials and Chilean Court ruling (2023).
Apple Inc. "Brain Computer Interface Human Interface Device input protocol" (2025) developer documentation.
Spherical Insights / Mordor Intelligence neurotechnology market reports (2025-2026).

Chapter 67-bis — The Ideal Legislative Stack

Organisation for the Prohibition of Chemical Weapons. Verification Annex to the Chemical Weapons Convention (1993).
International Atomic Energy Agency. Safeguards Agreement model (INFCIRC/153).
UNESCO Recommendation on the Ethics of Neurotechnology (2025) — implementation tools and Member State guidance.
Chile Law 21.383 (2021) implementing regulations; Brazil constitutional amendment implementing legislation drafts.
Council of Europe. Framework Convention on AI (2024) — model for binding cross-state instrument.

Model Law on the Sale of Goods (UNCITRAL) — model for drafting cross-jurisdictional substantive provisions.
Park doctrine: United States v. Park, 421 U.S. 658 (1975) — model for corporate-officer criminal liability provisions.

Chapter 70-bis — Implementation: Agencies

European Data Protection Board. Annual Reports; coordination mechanism documentation.
Organisation for the Prohibition of Chemical Weapons. Annual Reports; verification methodology publications.
U.S. Food and Drug Administration. 21 CFR 882 (Neurological Devices) regulatory framework.
National Institute for Standards and Technology (NIST). "AI Risk Management Framework" (2023) — model for technical-standard adoption.
Frischmann, B. & Verbeure, B. "Designing the Independent Regulator." American University Law Review (2021).

Chapter 76-bis — Geopolitical Impact

European Commission. "Brussels Effect" — Bradford, A. The Brussels Effect: How the European Union Rules the World. Oxford University Press (2020).
Allison, G. Destined for War: Can America and China Escape Thucydides's Trap? Houghton Mifflin (2017) — frame for US-China rivalry analysis.
Chinese Academy of Sciences. Brain Project public materials; PLA Military Brain Science programme open-source analyses.
Inter-American Commission on Human Rights. Resolutions and reports on neurotechnology (2022-2025).
Bradford, A. Digital Empires: The Global Battle to Regulate Technology. Oxford University Press (2023).

Chapter 78-bis — Synthetic Telepathy

Willett, F. R. et al. "High-performance brain-to-text communication via handwriting." Nature (2021) — Shenoy/Stanford imagined-speech decoding.
Metzger, S. L. et al. "A high-performance neuroprosthesis for speech decoding and avatar control." Nature (2023) — UCSF Chang group.
Rao, R. P. N. et al. "A Direct Brain-to-Brain Interface in Humans." PLOS One (2014) — first human-to-human BCI demonstration.
Boyden, E. S. "Optogenetics and the future of neuroscience." Nature Neuroscience (2015) — optogenetics substrate.
Legon, W. et al. "Transcranial focused ultrasound modulates the activity of primary somatosensory cortex in humans." Nature Neuroscience (2014).
Yuste, R. & Goering, S. "It's Time for Neuro-Rights." Horizons (2020) — foundational discussion of constraints.

Part XV-bis — Hypothetical Cases

GDPR Article 7(4) on freedom of consent in imbalanced relationships; Article 29 Working Party Opinion 02/2017 on employee consent.
EU AI Act Article 5 (prohibited practices) and Annex III (high-risk systems).
Brazilian LGPD (Lei 13.709/2018); Statute of the Child and Adolescent (Lei 8.069/1990).
French Code de la Consommation; Code Civil; loi Informatique et Libertés.
California Consumer Privacy Act, Civil Code § 1798.140 et seq.; MIND Act (proposed federal, 2025).
Vienna Convention on Diplomatic Relations (1961); Geneva Conventions (1949) and Additional Protocols (1977).
Tallinn Manual 2.0 on the International Law Applicable to Cyber Operations (NATO CCDCOE, 2017).

International Criminal Court Rome Statute (1998), Articles 7, 8.
Frye v. United States, 293 F. 1013 (D.C. Cir. 1923); Daubert v. Merrell Dow Pharmaceuticals, 509 U.S. 579 (1993) — evidentiary admissibility standards potentially applied to memory-modification cases.
Title VII of the Civil Rights Act of 1964 (US); Americans with Disabilities Act (US, 1990).

Section 5 — Citation Practice and Update Note

The book uses an in-text descriptive citation practice rather than footnote or endnote markers. Readers who wish to locate the source for a specific claim should consult the chapter-relevant section of this bibliography; in most cases the source is identifiable from the descriptive citation in the chapter text. References in this section reflect the state of the field as of mid-2026 and will require periodic updating; readers should consult the TI One Voice platform at one.witysk.org for current revisions.

Permission to reproduce portions of this book for educational, advocacy and legislative purposes is granted, with attribution. The book and its references are made available under the same open principles that animate the broader TI One Voice initiative.