

# From Ethics to Governance: Why Neurotechnology Requires Institutionally Actionable Concepts

Yong Ding 

Weifang University of Science and Technology, Weifang, China | [ddy2004@126.com](mailto:ddy2004@126.com)

Received: 20 February 2026 Revised: 3 March 2026 Accepted: 4 March 2026 Published: 15 March 2026

## Abstract

Neurotechnology governance confronts a paradox: the proliferation of ethical concepts — cognitive liberty, mental privacy, neurorights, mental self-determination — has expanded normative vocabulary without producing commensurate institutional capacity for enforcement. This article argues that the governance deficit stems not from insufficient ethical ambition but from a structural mismatch between the concepts scholars produce and the operational grammar institutions require. Regulatory authorities cannot process ethical aspirations; they can only act on concepts expressible as observable triggers, evidentiary substrates, competent authorities, and enforceable remedies. To diagnose this mismatch, the article develops the Institutional Actionability Assessment (IAA), a replicable diagnostic framework that evaluates any governance-relevant concept along three indicators — Verification Void, Authority Fit, and Trigger Executability — yielding a bounded composite metric, the IAC-Index. The IAA is applied to four dominant neuroethical concepts and one cross-technology transferability case, revealing that each concept fails on at least two dimensions of institutional readability. The article then demonstrates “downward translation” by compiling mental privacy into an Institutionally Actionable Proxy specifying trigger, evidence, authority, and remedy. Five design principles for constructing actionable proxies are articulated. A failure-mode typology maps diagnostic score patterns to institutional repair levers, closing the gap between diagnosis and prescription. The framework is deliberately architecture-portable: it evaluates concepts regardless of which regulatory stack is in place, making findings portable across jurisdictions and technology domains. The IAA equips policymakers, regulatory designers, and standards bodies with a replicable diagnostic tool for identifying and repairing institutional actionability failures in any emerging technology domain.

**Keywords** neurotechnology governance, neuroethics, cognitive liberty, mental privacy, institutionally actionable concepts, regulatory theory, actionability gap, enforcement readability, conceptual translation, emerging technology governance, IAC-Index, neurorights, downward translation, institutional readability

## 1 Introduction

### 1.1 Why Neurotechnology Now Requires Governance-Executable Concepts

Neurotechnology governance fails not because ethics lacks ambition, but because ethical concepts often lack *institutional actionability*: they cannot be parsed into measurable triggers, evidentiary substrates, competent authorities, and enforceable remedies. The neuroethics literature has experienced extraordinary conceptual productivity since the mid-2010s: cognitive liberty, mental privacy, neurorights, neural data sovereignty, and the right to mental self-determination have entered academic, legislative, and policy discourse at an accelerating rate (Ienca & Andorno, 2017); (Bublitz, 2013); (Farahany, 2023); (Yuste et al.,

2017). This conceptual proliferation has generated the appearance of governance progress. The reality is more troubling.

A bibliometric scan of Web of Science Core Collection records (search query: TS=(“neurorights” OR “cognitive liberty” OR “mental privacy” OR “neural data sovereignty”), retrieved 28 February 2026) reveals that publications containing these terms in their title, abstract, or keywords grew from fewer than 20 per year before 2017 to over 200 by 2025 (Table 1). Yet legislative conversion remains near-zero: Chile’s 2021 constitutional amendment (Art. 19 No. 1) stands as the sole binding national-level codification, and it remains institutionally inert — no implementing statute, no designated enforcement authority, no compliance test has been enacted as of March 2026 (Senate of Chile, 2021). The IEEE Neuroethics Framework (2024), the UNESCO recommendation on the ethics of neurotechnology (2023), and the OECD Responsible Innovation in Neurotechnology recommendation (2019) have proliferated governance *principles* without producing enforcement *syntax* (OECD, 2019); (UNESCO, 2023). What the field has produced, in institutional terms, is an expanding inventory of normative aspirations disconnected from any mechanism that would give them teeth.

This article is deliberately architecture-portable: it diagnoses the institutional readability of ethical concepts regardless of which regulatory stack is in place, making its findings applicable beyond any single jurisdiction’s governance design. Whether the regulatory architecture is the EU’s layered MDR-GDPR regime, the US patchwork of FDA and state privacy laws, or Chile’s constitutional model, the diagnostic question is the same: can the concepts embedded in governance discourse be parsed into the operational grammar that institutions require for action?

Table 1: Neuroethics governance publications in Web of Science Core Collection (TS query, retrieved 28 February 2026)

Year	2010	2013	2015	2017	2019	2021	2023	2024	2025
Annual count	4	8	12	19	38	72	145	189	216
Cumulative	18	33	52	82	148	294	582	771	987

Note. Query: TS=(“neurorights” OR “cognitive liberty” OR “mental privacy” OR “neural data sovereignty”). Database: Web of Science Core Collection. Timespan: 2005–2025. Duplicate removal: by DOI. Counts rounded to nearest integer from monthly WoS indices.

## 1.2 The Actionability Gap and the Translation Paradox

The accumulation of neuroethics concepts may be counterproductive for neurotechnology governance if those concepts remain institutionally inert: they create the appearance of normative coverage while foreclosing more modest but operable regulatory interventions. This is the *actionability gap* — the structural distance between what ethical discourse produces (aspirational concepts) and what institutions consume (enforceable triggers, evidentiary standards, designated authorities, and defined remedies). The gap is not incidental; it is produced by a disciplinary division of labour in which normative theorists and institutional designers operate with different grammars, different success criteria, and different audiences.

The concept of “translation” invoked here — downward translation from ethical aspiration to operational proxy — overlaps with established literatures but is analytically distinct from each. It differs from Actor–Network Theory’s notion of translation (Callon, 1986), which concerns the enrollment of allies across social worlds; this article’s translation concerns the *compilation* of concepts into enforceable institutional format. It engages the capabilities approach (Sen, 1999); (Nussbaum, 2011) as a precedent for operationalizing abstract rights into measurable functionings — but adds a specifically institutional dimension that capabilities theory leaves under-specified: who enforces, what triggers enforcement, and what

---

For a complementary architectural analysis of *where* the EU governance stack fails for neurotechnology — as opposed to this article’s focus on *why* the concepts themselves resist institutional processing — see Li (2026) in this issue.

evidence sustains it. It acknowledges the boundary objects framework (Star & Griesemer, 1989) as relevant to how concepts travel between social worlds (ethics → regulation), but moves beyond description to prescription: the task is not merely to observe how concepts are “translated” across communities but to engineer proxies that institutions can actually process.

### 1.3 Contribution: Institutional Actionability Assessment (IAA)

This article makes three core contributions to the political and institutional theory of neurotechnology governance:

- **First, it formalises the actionability gap through three propositions** — the Translation Paradox, the Enforcement Readability Constraint, and the Downward Translation Imperative — each with analytical sketches and falsification conditions.
- **Second, it develops the *Institutional Actionability Assessment (IAA)***, a replicable diagnostic framework scoring any governance-relevant concept on three indicators (Verification Void, Authority Fit, Trigger Executability) and yielding a bounded composite metric, the IAC-Index.
- **Third, it demonstrates “downward translation”** by compiling mental privacy and cognitive liberty into *Institutionally Actionable Proxies (IAPs)* and articulates five design principles for constructing such proxies. A failure-mode typology maps diagnostic patterns to institutional repair levers, closing the diagnosis-to-prescription loop.

The framework is designed as a *reusable diagnostic tool* for governance scholars, policymakers, and regulatory designers working on any emerging technology where ethical concepts outpace institutional capacity — including but not limited to neurotechnology, affective computing, digital therapeutics, and synthetic biology.

## 2 Background: The Rise of Neuroethics as a Governance Discourse

### 2.1 From Laboratory Ethics to Public Policy Discourse

Neuroethics emerged as a recognisable field in the early 2000s, initially anchored in the ethical oversight of neuroscience research — questions of informed consent for brain imaging, the moral status of cognitive enhancement, and the implications of neurological determinism for legal responsibility (Illes & Raffin, 2006); (Roskies, 2002). By the mid-2010s, the field had undergone a consequential transformation: it moved from laboratory-level bioethics to public-policy-level governance discourse. The catalyst was a convergence of technological maturation (consumer EEG devices, invasive BCI clinical trials, direct-to-consumer neurostimulation), political attention (the BRAIN Initiative in the US, the Human Brain Project in the EU), and normative entrepreneurship by scholars who explicitly sought to link neuroethical concepts to rights frameworks and legislative action (Ienca & Andorno, 2017); (Yuste et al., 2017).

This transformation was productive: it brought neurotechnology into governance discussions that had previously been confined to clinical and research ethics committees. But it also introduced a structural tension. The concepts that proved most successful in capturing public and legislative attention — cognitive liberty, mental privacy, neurorights — were precisely those formulated at the highest level of abstraction, in the register of fundamental rights and constitutional principles. They were designed to *declare what ought to be protected*, not to *specify how institutions should protect it*. The gap between declaration and specification is the terrain this article occupies.

## 2.2 Key Concepts and Their Institutional Reception

Four concepts dominate the neuroethics governance lexicon, each with a distinct genealogy and a distinctive relationship to institutional operationalization. These four were selected as the highest-frequency rights-cluster in OECD, UNESCO, and Council of Europe neurotechnology policy texts and as the most-cited governance-relevant concepts in the neuroethics literature (2017–2025), ensuring that the diagnostic exercise engages the concepts with the greatest institutional salience rather than a convenience sample:

**Cognitive liberty** was articulated by (Bublitz (2013, 2015)) as a negative freedom: the right to be free from non-consensual interference with one’s cognitive processes, encompassing both the freedom to use neurotechnology and the freedom to refuse its use. The concept draws on liberal autonomy theory and locates itself in the register of fundamental rights. Its institutional reception has been largely symbolic: no jurisdiction has enacted enforceable cognitive liberty protections with specified triggers, evidence standards, or designated enforcement authorities.

**Mental privacy** was developed most influentially by (Farahany (2023)) and (Ienca and Andorno (2017)) as the right to control access to and disclosure of information derived from one’s neural activity. The concept is intuitively compelling and has generated substantial academic traction, but its institutional specification remains contested: where does GDPR-standard data privacy end and “mental” privacy begin? What specific data processing practices constitute a violation? How does a data protection authority distinguish a mental privacy breach from a routine data protection infringement?

**Neurorights and neural data sovereignty** encompass a family of claims about the ownership, control, and governance of brain-derived data (Yuste et al., 2017). Chile’s constitutional amendment (2021) is the most prominent legislative output, but it exemplifies the paradox this article diagnoses: the constitutional provision exists, but no implementing statute translates it into enforceable institutional action. The concept of “sovereignty” over neural data imports metaphors from indigenous data sovereignty and data localization debates without specifying which institutions can operationalize sovereignty claims.

**The right to mental self-determination** operates as a comprehensive covering concept: the right to determine one’s own mental states, cognitive development, and neural identity without external interference (Ienca, 2021). It is the most encompassing and the least institutionally specified of the four concepts. What prohibitions does it impose on neurotechnology developers? What enforcement mechanism gives it teeth? These questions remain substantially unanswered.

## 2.3 The Partial Successes and Structural Limits of Neuroethics Governance Claims

The neuroethics governance discourse has achieved genuine partial successes. It has placed neurotechnology on legislative agendas in Chile, Spain, and at the EU level. It has influenced the OECD and UNESCO to develop neurotechnology-specific principles. It has created a shared vocabulary through which scholars, policymakers, and civil society organizations discuss brain-related governance concerns. At the subnational level, Colorado (HB 24-1058, 2024) has extended its Privacy Act to cover biological and neural data (Colorado General Assembly, 2024), and Minnesota (SF 1240, 2025–2026) has introduced neurodata-specific privacy provisions (Minnesota Legislature, 2025). These represent incremental operationalisation within existing data privacy architectures; they raise AF scores for mental privacy in the US context but do not create neurotechnology-specific enforcement authorities and leave cognitive liberty entirely unaddressed. These contributions are real and should not be minimized.

The structural limit is equally real: the discourse has produced *normative inflation* without proportional *institutional conversion*. The number of concepts, principles, and rights-claims has grown exponentially; the number of enforceable institutional mechanisms has grown barely at all. This pattern is not unique to neuroethics — it recurs wherever normative discourse runs ahead of institutional design (Majone, 1997); (Black, 2002) — but its consequences for neurotechnology governance are particularly acute because the

technology is advancing rapidly, consumer deployment is expanding, and the window for effective institutional design may be closing.

### 3 The Actionability Gap: A Theoretical Account

#### 3.1 Proposition 1: The Translation Paradox

**Proposition 1 (Translation Paradox).** *Normative concept proliferation can increase governance risk by creating the appearance of progress while leaving institutions without executable governance functions.*

*Analytical sketch.* Concept inflation occupies political bandwidth, crowds out operationalisable alternatives, and gives legislators a false confidence that “rights exist” for the new technology. When Chile enshrined neurorights in its constitution, it signalled normative closure — the problem of brain-related rights was “solved” at the constitutional level — while the institutional machinery for enforcement remained entirely absent. The political capital invested in the constitutional provisions reduced the urgency (and the political feasibility) of enacting the implementing statutes that would be necessary for institutional action. Concept proliferation thus produces a governance mirage: the more concepts are declared, the more governance *appears* to advance, and the less pressure is exerted for the operational work of institutional design.

*Falsification condition.* If a jurisdiction converts an ethical concept into enforceable statute within 24 months of its academic formulation — with specified triggers, designated authorities, and defined remedies — Proposition 1 does not hold for that concept-jurisdiction pair.

#### 3.2 Proposition 2: The Enforcement Readability Constraint

**Proposition 2 (Enforcement Readability Constraint).** *Regulatory authorities can only act on concepts that can be expressed as: (i) observable trigger, (ii) evidentiary substrate, (iii) competent authority, (iv) enforceable remedy.*

*Analytical sketch.* Administrative procedure theory requires each of these four elements for lawful institutional action (Baldwin, Cave, & Lodge, 2012); (Majone, 1997). A regulatory authority cannot initiate an enforcement action without a triggering event that it can observe or receive reports of; it cannot sustain that action without evidence that meets its procedural standards; it cannot act at all unless it possesses the statutory mandate and operational capacity to do so; and it cannot complete the action without a defined remedy (sanction, injunction, redesign order, recall) that is proportionate and legally available. Missing *any one* of these four elements renders a concept unprocessable by the institutional machinery of enforcement. The concept may be normatively compelling, philosophically rigorous, and politically popular — but it remains institutionally inert.

*Falsification condition.* Soft-law platforms (recommendations, principles, multi-stakeholder frameworks) can operate on vaguer concepts — they do not require legally codified triggers or statutory remedies. Proposition 2 applies strictly to *binding enforcement actions* within formal regulatory architectures. Institutions operating in advisory, deliberative, or standard-setting modes can process more abstract concepts, though even these institutions benefit from operational specificity.

#### 3.3 Why These Flaws Are Structural, Not Incidental

The actionability gap is not the product of insufficient effort by neuroethicists, nor of regulatory negligence, nor of political indifference. It is a structural consequence of the disciplinary division of labour between normative theory and institutional design. Normative theorists are trained to articulate what

ought to be protected, to defend rights–claims, and to expose governance deficits. Institutional designers — regulatory lawyers, administrative governance scholars, compliance engineers — are trained to build mechanisms that translate normative commitments into enforceable action. These two communities operate with different grammars, different success metrics, and different audiences. The neuroethics discourse has been spectacularly successful by the standards of normative theory (concept recognition, constitutional codification, international attention); it has been largely unsuccessful by the standards of institutional design (enforcement readiness, compliance testability, remedy specification).

This structural diagnosis locates the problem not in any single actor’s failure but in a *systemic design flaw*: the governance pipeline lacks a translation module that would compile ethical concepts into institutional format. The IAA framework developed in the next section provides that module.

## 4 Method: Institutional Actionability Assessment

The IAA is a diagnostic framework that evaluates any governance–relevant concept’s readiness for institutional processing. Its unit of analysis is the **triple: concept × institution × jurisdiction**. A concept that is “inert” in one institutional–jurisdictional context may be actionable in another: cognitive liberty, for example, scores differently when assessed against EU market surveillance than against the Chilean constitutional court. All scores reported in this article are therefore context–bound; the IAA does not treat actionability as an intrinsic property of concepts but as a relational property of concept–institution–jurisdiction triples.

### 4.1 Indicators and Codebook (VV, AF, TE)

Three indicators operationalise the actionability gap. Each is scored on a 0–3 scale where 0 = fully actionable and 3 = fully inert.

Table 2: IAA Codebook: Indicators, Scoring Rules, Evidence, and Target Institutions

Indicator	Definition	Scoring (0–3)	Evidence Base	Target Institution
<b>Verification Void (VV)</b>	Can violation be detected without reading mental content?	0 = fully externally verifiable; 1 = partially verifiable from artefacts; 2 = requires inference about mental states; 3 = inherently unobservable	Device telemetry logs, API records, product specifications, audit reports	Market surveillance authority, DPA
<b>Authority Fit (AF)</b>	Is there a clear competent authority with mandate and capacity?	0 = named authority with budget and mandate; 1 = authority exists, mandate ambiguous; 2 = fragmented across multiple authorities; 3 = orphan (no authority owns it)	Statutory mandates, organisational charts, enforcement budgets, published guidance	Any regulatory or standard-setting body
<b>Trigger Executability (TE)</b>	Can the trigger be operationalised into inspectable thresholds?	0 = binary, inspectable; 1 = threshold-based, measurable; 2 = requires contextual judgement; 3 = inherently subjective	Compliance test protocols, inspection checklists, technical standards	Notified bodies, market surveillance, labour inspectorates

The Authority Fit indicator warrants additional theoretical specification. Neurotechnology governance inherently involves polycentric authority structures — multiple overlapping regulators operating at different jurisdictional levels with partial mandates (Ostrom, 2010). AF captures the degree to which this polycentricity produces institutional clarity (AF = 0, one clear lead authority) or institutional orphanhood (AF = 3, no authority claims the concept). The AF indicator thus provides an institutional diagnosis that

connects directly to (Ostrom’s) polycentric governance framework: high AF scores identify concepts that fall into the interstitial spaces between overlapping authorities.

**Why remedy is not indexed quantitatively.** Proposition 2’s four-element syntax includes *remedy*, but the IAC-Index quantifies only VV, AF, and TE. This asymmetry is deliberate: remedy is *downstream* of the other three elements. A remedy can only attach *after* a trigger has been detected (TE), evidence has been gathered (VV), and a competent authority has been identified (AF). Absent any one of these preconditions, remedy specification is premature. Remedy is therefore treated as a *qualitative diagnostic* in the worked example (Section 6) rather than a quantitative indicator — it is assessed in the Concept-to-Code compilation but does not enter the composite formula. This design choice prevents the index from double-counting institutional failures (a concept that lacks an authority also, by definition, lacks a remedy pathway through that authority).

## 4.2 The IAC-Index (Composite Metric)

The IAC-Index provides a bounded composite measure of institutional actionability:

$$\text{IAC-Index} = 1 - \frac{w_{\text{VV}} \cdot \text{VV} + w_{\text{AF}} \cdot \text{AF} + w_{\text{TE}} \cdot \text{TE}}{3 \cdot (w_{\text{VV}} + w_{\text{AF}} + w_{\text{TE}})}$$

Default: equal weights ( $w = 1$ ), yielding IAC-Index  $\in [0, 1]$ . Three reporting bands are defined:  $\geq 0.67 =$  **Actionable** (institution-ready; concept can be processed by existing regulatory machinery with minimal modification);  $0.33\text{--}0.67 =$  **Partially actionable** (requires targeted translation work on the failing indicator(s));  $< 0.33 =$  **Inert** (aspirational only; full downward translation via IAP design is required).

The IAC-Index is presented as a *replication-ready metric for future empirical studies*: product mapping (scoring consumer neurodevices against these indicators), complaint pathway tracing (tracking institutional responses to alleged violations), and enforcement correspondence analysis (examining how regulators process neuroethics-adjacent complaints). The current application is illustrative; the metric’s value lies in its replicability. We offer the IAA as an open methodology: the Box 1 codebook, the three indicators, and the IAC-Index formula are designed for direct adoption by empirical scholars evaluating governance readiness in any technology domain. Researchers may adapt the weighting scheme, extend the indicator set, or embed the protocol within larger comparative designs; the baseline architecture is intentionally minimal to maximise portability.

### Box 1: The IAA Diagnostic Checklist

**Part A — Coding Protocol.** *Evidence priority*: primary legal text > official guidance > peer-reviewed analysis > grey literature; score only on the basis of the highest-priority evidence available. *Missing information*: if no evidence is available for a cell, assign score = 3 (worst case) and flag “[insufficient evidence].” *Inter-coder reliability*: scoring is designed for multi-coder application; recommended minimum protocol: two independent coders score all cells; compute Cohen’s  $\kappa$  for each indicator;  $\kappa \geq 0.60 =$  acceptable agreement; disagreements resolved by structured discussion with documented rationale; unresolved items default to the higher (less actionable) score; pilot testing with 2+ coders on the four concepts in this article is recommended before field deployment.

**Part B — Diagnostic Steps.** (i) Identify the **normative invariant** of the concept (the irreducible semantic core that any actionable proxy must preserve). (ii) For the target institution-jurisdiction pair, score VV, AF, and TE using the codebook (Table 1). (iii) Compute IAC-Index and assign to reporting band. (iv) If IAC-Index  $< 0.67$ : identify the dominant failure mode (VV-dominant, AF-dominant, TE-dominant, or triple failure). (v) If IAC-Index  $< 0.33$ : proceed to downward translation (Part C).

**Part C — Compilation (Downward Translation).** (i) State the normative invariant explicitly. (ii) Specify: trigger (what measurable event?), evidence (what artefacts?), authority (who acts and why?), remedy (what enforceable action?), boundary conditions (when does the IAP not apply?). (iii) Verify fidelity: does the proxy preserve the normative invariant? If not, revise. (iv) Apply the five design principles (Section 6.2).

Table 3: IAA Empirical Coding Handbook

Indicator	Score & Meaning	Acceptable Evidence & Examples
Verification (VV)	0: Fully verifiable	<i>Evidence:</i> API logs, source code, data flow audits. <i>Example:</i> Local vs. cloud storage location.
	1: Partially verifiable	<i>Evidence:</i> Privacy policies, developer terms. <i>Example:</i> Declared data sharing practices.
	2: Inference required	<i>Evidence:</i> Post-hoc analysis of model outputs. <i>Example:</i> Ascertaining emotional inference from EEG.
	3: Inherently unobservable	<i>Evidence:</i> Self-reporting only. <i>Example:</i> Internal cognitive disruption.
Authority Fit (AF)	0: Dedicated authority	<i>Evidence:</i> Statutory mandate, specific enforcement budget. <i>Example:</i> FDA mandate over medical BCI devices.
	1: Ambiguous mandate	<i>Evidence:</i> Regulator guidance documents. <i>Example:</i> DPA applying GDPR to neurodata.
	2: Fragmented authority	<i>Evidence:</i> Overlapping jurisdictions without lead agency. <i>Example:</i> Consumer safety vs. data protection.
	3: Institutional orphan	<i>Evidence:</i> No regulatory body claims jurisdiction. <i>Example:</i> “Mental self-determination” enforcement.
Trigger Executability (TE)	0: Binary inspectable	<i>Evidence:</i> Technical compliance protocols. <i>Example:</i> “Data left the device” (Yes/No).
	1: Threshold measurable	<i>Evidence:</i> Quantitative limits in standards. <i>Example:</i> Maximum allowable inference depth.
	2: Contextual judgement	<i>Evidence:</i> Balancing tests. <i>Example:</i> Proportionality of workplace neural monitoring.
	3: Inherently subjective	<i>Evidence:</i> Philosophical or constitutional claims. <i>Example:</i> “Interference with cognitive liberty”.

## 5 Stress-Testing the Neuroethics Lexicon

For each of the four neuroethical concepts, the normative invariant is identified first, then VV, AF, and TE are scored against the EU regulatory context (GDPR + MDR + national market surveillance). The assessment context is the *EU post-AI-Act institutional landscape as of March 2026*; scores in US or Chilean contexts would differ, demonstrating the relational nature of actionability.

### 5.1 Cognitive Liberty in Clinical and Consumer Scenarios

**Normative invariant:** freedom from non-consensual interference with cognitive processes.

**VV = 3.** What counts as “interference with cognitive processes” in a way that a market surveillance authority can observe from device telemetry, API logs, or product specifications? Consumer EEG headbands acquire neural signals and provide neurofeedback that can condition cognitive patterns — but whether this constitutes “interference” with “cognitive processes” requires interpretation of mental states that no external artefact can supply. A consumer who uses a focus-enhancement headband is voluntarily engaging with the device; determining whether a subsequent employer mandate to use such a device constitutes “non-consensual interference” requires judgements about coercion that are contextual, not inspectable.

*Scenario:* An employer deploys workplace neural monitoring (e.g., SmartCap-style EEG for fatigue

detection) without meaningful opt-out. Does this violate cognitive liberty? The concept provides no threshold for “non-consensual” — no trigger that would activate an enforcement response.

**AF = 3.** No EU authority has a mandate that encompasses “cognitive liberty” as an enforceable standard. DPAs enforce data protection, not cognitive freedom. Labour inspectorates enforce workplace safety, not cognitive autonomy. The EDPS has published analytical reports on neurodata but has no enforcement mandate over cognitive liberty claims. The concept is an institutional orphan: it commands normative authority but no bureaucratic budget.

**TE = 3.** The trigger condition — “non-consensual interference with cognitive processes” — is inherently subjective. It specifies neither the act that constitutes interference nor the threshold at which interference becomes actionable. No compliance test protocol could operationalise it.

**IAC-Index** =  $1 - \frac{3+3+3}{9} = 0.00$ . *Fully inert.*

## 5.2 Mental Privacy: Where Biometric Ends and Neural Begins

**Normative invariant:** control over disclosure of information derived from neural activity.

**VV = 2.** Mental privacy is partially verifiable: a DPA can inspect whether a consumer EEG device collects neural data and transmits it to cloud servers (this is externally observable from data flow analysis and privacy policy audits). But the *mental* dimension — whether the processed data reveals information about the user’s mental states, cognitive patterns, or emotional responses — requires inference about what the data *can reveal*, not merely what it *does reveal*. The boundary between standard biometric data processing and “mental privacy” violation is computationally invisible: the same EEG signal that supports a meditation score *could* support cognitive profiling, and no external audit can determine which inferences are actually being drawn.

*Scenario:* A consumer EEG meditation app processes raw EEG telemetry in the cloud. The app’s stated purpose is “relaxation scoring.” The same data could, in principle, support attention-state classification, emotional valence inference, or implicit attitude detection (Martinovic et al., 2012). Has mental privacy been violated? The answer depends on what inferences are *actually drawn* — information that is opaque to external audit without compulsory process.

**AF = 2.** Data protection authorities have a partial mandate: GDPR applies to neural data processing, and the EDPS TechDispatch #1/2024 suggests that neurodata is likely special-category data under Article 9 (EDPS, 2024). But DPAs lack sector-specific auditing capacity for neural data. Market surveillance authorities lack data governance mandates. The result is a fragmented authority landscape where no single regulator possesses the combination of neural-data expertise and enforcement mandate needed for comprehensive mental privacy protection.

**TE = 2.** Partial trigger executability exists: “data left the device” is binary and inspectable; “data was processed in the cloud” is verifiable from network traffic analysis. But the core trigger — “information *derived from neural activity* was disclosed without adequate consent” — requires contextual judgement about what constitutes “adequate consent” for neural data and what processing operations constitute “disclosure” versus legitimate service provision.

**IAC-Index** =  $1 - \frac{2+2+2}{9} = 0.33$ . *Borderline: at the lower edge of “partially actionable.”*

## 5.3 Neural Data Sovereignty and Jurisdictional Fragmentation

**Normative invariant:** collective or individual right to control the conditions under which neural data is accessed, used, and governed.

**VV = 1.** Data flows are substantially observable: where neural data is stored (local vs. cloud), which entities access it (first-party vs. third-party SDKs), and across which jurisdictions it moves are all verifiable through technical audit and privacy policy analysis.

**AF = 3.** “Sovereignty” is a jurisdictional concept that maps poorly onto existing regulatory authority structures. DPAs enforce data protection based on legal bases (consent, legitimate interest, public interest) — not sovereignty claims. No authority has a mandate to enforce “neural data sovereignty” as a distinct governance standard, and the concept draws competing jurisdictional analogies (indigenous data sovereignty, data localization, digital sovereignty) that different institutions would interpret in incompatible ways.

**TE = 2.** Some elements are threshold-capable: data localization requirements (“neural data must be stored within the jurisdiction”) can be technically enforced. But the broader sovereignty claim — “the individual or community has ultimate control over the conditions of neural data governance” — is not reducible to inspectable thresholds. It describes a state of affairs, not a testable compliance condition.

**IAC-Index** =  $1 - \frac{1+3+2}{9} = 0.33$ . *Borderline: at the lower edge of “partially actionable.”*

#### 5.4 Mental Self-Determination as Constitutional Aspiration

**Normative invariant:** the right to determine one’s own mental states, cognitive development, and neural identity without external interference.

**VV = 3.** “Determining one’s own mental states” is a first-person phenomenological claim that no institutional audit can verify or falsify from external artefacts. Even with full access to device telemetry and data flows, a market surveillance authority cannot determine whether a user’s “mental self-determination” has been compromised.

**AF = 3.** No institutional authority has a mandate over “mental self-determination.” The concept operates at a constitutional-philosophical level that is structurally resistant to administrative assignment.

**TE = 3.** The concept imposes no identifiable prohibitions or requirements on technology developers. No compliance test, inspection protocol, or technical standard could operationalise “respect for mental self-determination.”

**IAC-Index** =  $1 - \frac{3+3+3}{9} = 0.00$ . *Fully inert.*

#### 5.5 Pattern Analysis: What the Diagnostics Reveal

Table 4: Concept Stress-Test: IAA Scores in the EU Institutional Context

Concept	VV	AF	TE	IAC-Index	Failure Mode	Repair Lever
Cognitive liberty	3	3	3	0.00	Triple failure	Full IAP design
Mental privacy	2	2	2	0.33	Balanced partial	Threshold codification + authority clarification
Neural data sovereignty	1	3	2	0.33	AF-dominant	Mandate assignment or inter-authority MoU
Mental self-determination	3	3	3	0.00	Triple failure	Full IAP design
<i>Emotional privacy (affective computing)</i>	2	2	1	0.44	<i>VV/AF partial</i>	<i>Standardise emotion-inference auditing</i>

**Cross-technology transferability test.** The final row of Table 2 applies the IAA to a non-neurotech concept: “emotional privacy” in the context of affective computing. The concept of a “right to emotional privacy” — the right not to have one’s emotional states inferred and used without consent — faces similar

but not identical actionability challenges.  $VV = 2$  (emotion inference from facial analysis or voice prosody is partially externally auditable, but the accuracy and meaning of inferred emotional states remain contested).  $AF = 2$  (DPAs and AI Act enforcement authorities have partial but fragmented mandates for emotion recognition systems).  $TE = 1$  (the EU AI Act’s prohibition on emotion recognition in certain contexts provides a partially operational trigger, making this concept more trigger-executable than most neuroethics concepts). This row demonstrates that the IAA framework is not a neurotech-specific rhetorical device; it diagnoses institutional actionability gaps across technology domains.

Three patterns emerge from the diagnostic results:

*First*, all four neuroethics concepts score below the threshold for institutional actionability ( $< 0.67$ ), and two (cognitive liberty, mental self-determination) are fully inert ( $= 0.00$ ). This is not an artefact of overly stringent scoring; it reflects the structural absence of enforceable triggers, designated authorities, and inspectable thresholds for these concepts.

*Second*, the concepts that score highest (mental privacy, neural data sovereignty) are those that overlap with *existing* regulatory frameworks (GDPR, data localization). Their partial actionability derives not from their own design but from the pre-existing institutional infrastructure of data protection. Concepts that lack this institutional anchor (cognitive liberty, mental self-determination) are entirely inert.

*Third*, the cross-technology test (emotional privacy) scores higher than most neuroethics concepts, largely because the EU AI Act provides a partially operational trigger for emotion recognition — specifically, Article 5(1)(f) prohibits emotion-inference systems in certain high-risk contexts, supplying precisely the kind of observable threshold that the neuroethics lexicon lacks. This suggests that *legislative specificity* — even imperfect specificity — substantially improves actionability, and that the neuroethics discourse’s preference for abstract, comprehensive rights-claims may be counterproductive for institutional conversion. The pattern also illuminates why Authority Fit varies so sharply across concepts: concepts that map onto polycentric governance structures (cf. (Ostrom, 2010)) — where multiple authorities hold partial but complementary mandates — achieve higher AF scores than concepts that presuppose a single, purpose-built regulator that does not yet exist. This polycentricity dividend explains the relative success of data-anchored concepts (mental privacy, neural data sovereignty) and the failure of philosophically comprehensive ones (cognitive liberty, mental self-determination).

## 5.6 Robustness Check: Weight Sensitivity

The IAC-Index as presented uses equal weights ( $w_{VV} = w_{AF} = w_{TE} = 1$ ). A natural objection is that Authority Fit might be a necessary precondition — without a competent authority, neither verification infrastructure nor trigger specification matters — warranting higher weight. Table 5 reports IAC-Index scores under three weighting schemes.

Table 5: IAC-Index Sensitivity to Alternative Weighting Schemes

Concept	Equal weights		AF-priority ( $w_{AF} = 2$ )		TE-priority ( $w_{TE} = 2$ )	
	Score	Band	Score	Band	Score	Band
Cognitive liberty	0.00	Inert	0.00	Inert	0.00	Inert
Mental privacy	0.33	Partial	0.33	Partial	0.33	Partial
Neural data sov.	0.33	Partial	0.28	Inert	0.33	Partial
Mental self-det.	0.00	Inert	0.00	Inert	0.00	Inert
Emotional privacy	0.44	Partial	0.42	Partial	0.47	Partial

*Note.* AF-priority:  $w_{AF} = 2, w_{VV} = w_{TE} = 1$ . TE-priority:  $w_{TE} = 2, w_{VV} = w_{AF} = 1$ . Bands: Actionable  $\geq 0.67$ ; Partial 0.33–0.66; Inert  $< 0.33$ .

The key finding is that the *rank ordering is invariant* to reasonable weight perturbations. Cognitive

liberty and mental self-determination remain fully inert under all schemes; emotional privacy remains the highest-scoring concept. The only classification shift occurs for neural data sovereignty, which drops from “Partial” to “Inert” under AF-priority weighting — a theoretically coherent result, since the concept’s primary vulnerability is its lack of institutional home (AF = 3). This confirms that equal weights are a defensible default and that the framework’s diagnostic conclusions are robust.

## 5.7 Architecture Portability: EU vs. US Micro-Comparison

The IAA framework claims architecture portability: scores should vary across institutional contexts, demonstrating that actionability is a relational property of the concept × institution × jurisdiction triple. Table 6 presents a minimal cross-jurisdictional comparison.

Table 6: Cross-Jurisdiction IAA Scores: EU vs. US Context (March 2026)

Concept	EU (GDPR + MDR + AI Act)			US (FTC + state patchwork)		
	VV/AF/TE	IAC	Band	VV/AF/TE	IAC	Band
Mental privacy	2/2/2	0.33	Partial	2/2/2	0.33	Partial
Cognitive liberty	3/3/3	0.00	Inert	3/3/3	0.00	Inert

*Note.* US context: FTC unfairness authority provides partial data-practice mandate (AF = 2 for mental privacy); no federal cognitive liberty standard exists (AF = 3). Colorado HB 24-1058 extends state privacy to neural data but does not establish neurotechnology-specific enforcement authority.

Two findings warrant emphasis. First, cognitive liberty’s triple failure (3/3/3) is not an artefact of EU-specific regulation but a structural property of the concept’s current formulation — the scores are identical across jurisdictions. Second, mental privacy converges at 0.33 in both jurisdictions but for different institutional reasons: the EU’s partial score derives from GDPR infrastructure, while the US score derives from FTC unfairness authority and state-level neural data extensions (Colorado, Minnesota). The failure modes are jurisdictionally specific even when composite scores converge, underscoring that the IAA diagnoses institutional mechanisms, not merely aggregate readiness.

## 6 Engineering Actionability: The Downward Translation Framework

### 6.1 Proposition 3: The Downward Translation Imperative

**Proposition 3 (Downward Translation Imperative).** *High-level neuroethical concepts must be “compiled downward” into minimal operational proxies that preserve normative intent while enabling inspection and enforcement.*

*Analytical sketch.* The gap between aspiration and enforcement is bridged not by refining the aspiration but by engineering a proxy that institutions can process without normative loss. The analogy is to software compilation: a high-level programming language expresses intent in human-readable form; a compiler translates that intent into machine-executable code. The translation necessarily involves loss of expressiveness, but the compiled output is what actually runs. Similarly, ethical concepts express normative intent in philosopher-readable form; downward translation compiles them into institution-executable format. The resulting *Institutionally Actionable Proxy* (IAP) is less expressive than the original concept but is the version that governance machinery can actually process. The analogy is instructive but imperfect: software compilation preserves functional equivalence between source and object code, whereas downward translation necessarily involves semantic narrowing — the IAP captures a *subset* of the concept’s normative scope. The normative invariant check (Box 1, Part C, Step 3) is the mechanism that constrains how much narrowing is tolerable.

*Falsification condition.* Constitutional rights that function as absolute prohibitions — the ban on torture, the prohibition of slavery — need no downward translation. They are self-executing: the prohibited act is categorically identifiable, and no institutional discretion is required for enforcement. Proposition 3 applies to concepts that require *contextual institutional judgement* for enforcement, which includes all four neuroethics concepts examined in this article.

## 6.2 Design Principles for Institutionally Actionable Proxies (IAPs)

Five design principles govern the construction of IAPs. These principles are **prescriptive** — they tell concept-builders *how* to construct an IAC. The VV/AF/TE indicators (Section 4.1) are **diagnostic** — they tell evaluators *whether* an existing concept is actionable. The two are complementary: the principles specify the design standard; the indicators measure compliance with that standard.

1. **Verifiability Principle.** Conditions of satisfaction must be observable or inferable through institutional procedures — device telemetry, data flow audits, compliance documentation — without requiring access to the subject’s mental states.
2. **Institutional Home Principle.** Every concept must specify the authority-type empowered to act on it: DPA, market surveillance, notified body, labour inspectorate, or ethics committee. A concept without an institutional home is a concept without enforcement.
3. **Trigger Precision Principle.** The concept must specify the threshold or event that activates institutional response. “Violation of cognitive liberty” is not a trigger; “transmission of raw EEG telemetry to a cloud server without documented user consent to neural data processing” is a trigger.
4. **Remedy Adequacy Principle.** The proposed institutional response must be proportionate and operationally feasible. The remedy must fit within the designated authority’s existing enforcement toolkit or must specify the legislative change needed to create the remedy.
5. **Falsifiability Principle.** The concept must specify conditions under which it does *not* apply. A concept that purports to apply universally and without exception is a concept that applies nowhere in practice (Black, 2002).

## 6.3 Concept-to-Code Matrix and a Worked Example (Mental Privacy)

### Worked example: compiling mental privacy into an IAP.

*Aspirational (inert):* “Protecting the innermost thoughts of the user from external intrusion.”

*Normative invariant to preserve:* control over disclosure of information derived from neural activity.

*Actionable proxy (IAP):* “Mandatory local-only processing of raw EEG telemetry with a hardcoded maximum inference depth, auditable via predefined market surveillance API endpoints.”

The IAP is specified across the four-element institutional syntax. The *trigger* activates when raw neural signal data (EEG, fNIRS, or equivalent) is transmitted from the acquisition device to any server not physically co-located with the device, or when the device’s processing pipeline generates inferences beyond Level-1 (signal quality, basic spectral features) without documented user consent specifying the inference type. The *evidentiary substrate* comprises data flow audits of device-to-server communication, API call logs, inference model documentation (model cards), DPIAs where required, and product specification sheets documenting inference depth. The *authority* is dual: the national DPA enforces data processing compliance, while the market surveillance authority enforces product compliance with inference-depth labelling, both acting through existing GDPR Article 58 powers and market surveillance regulation. The *remedy* comprises mandatory local-only processing orders for raw neural data, inference-depth caps enforced through technical standards (maximum permissible inference level without enhanced consent), GDPR Article 83 administrative fines for non-compliant data transfers, and post-market monitoring duties requiring annual

inference–depth audits. The IAP does not apply to clinical neurotechnology classified as medical devices under MDR (governed by a separate regulatory pathway), academic research conducted under institutional ethics committee approval, or devices that process all neural data locally and generate no inferences beyond Level-1.

Table 7: Concept-to-Code Matrix: Downward Translation of Mental Privacy

Translation Level	Institutional Syntax	Formulation	Target Institution
Philosophical	Normative aspiration	“The right to mental privacy”	Academic discourse
Constitutional	Fundamental right	“Protection of the integrity of mental processes” (cf. Chile Art. 19 No. 1)	Constitutional court
Statutory	Legal obligation	“Processing of neural data requires explicit opt-in consent specifying inference types; raw neural data must be processed locally unless user consents to transfer”	Legislature → DPA
Technical proxy	Compliance test	“API endpoint returns TRUE if raw EEG leaves device; inference-depth manifest declares all inference levels beyond L-1”	Market surveillance + notified body

#### 6.4 Second Worked Example: Compiling Cognitive Liberty (Triple Failure)

The mental privacy example (Table 3) demonstrates downward translation for a *borderline* concept (IAC = 0.33). But the framework must also handle *triple-failure* cases (IAC = 0.00), where no existing institutional infrastructure provides a starting point. Cognitive liberty — the right to freedom from non-consensual interference with cognitive processes — presents this harder case. A crucial caveat is necessary: the triple-failure score is a *structural* diagnosis of the concept’s current institutional readiness, not a judgement on its normative worth. Cognitive liberty may well be the most important ethical principle in the neurotechnology domain; the IAA merely shows that, in its current philosophical form, it cannot be processed by existing regulatory machinery without active downward compilation.

**Normative invariant to preserve:** freedom from non-consensual interference with one’s cognitive processes.

*Aspirational (inert):* “Every person has the right to cognitive liberty: the freedom to control their own cognitive processes and to be free from non-consensual interference with their mental states.”

*Actionable proxy (IAP):* “Employers, educational institutions, and public service providers may not require, coerce, or incentivise the use of neurotechnology devices that record or modulate neural activity as a condition of employment, educational participation, or service access, unless authorised by sector-specific statutory exemption with a documented legitimate aim and independent ethics committee approval.”

The IAP is specified across the four-element institutional syntax. The *trigger* activates when an entity in an asymmetric-power relationship — employer, educational institution, or public service provider — conditions access to employment, education, or services on the use of a neural recording or neurostimulation device, or deploys such devices without documented opt-out mechanisms that do not penalise refusal.

The *evidentiary substrate* comprises employment contracts and workplace policies referencing neural monitoring, educational programme requirements specifying neurotechnology use, public service access procedures conditional on neural device use, opt-out documentation with records of consequences for refusal, and device procurement and deployment logs.

The *authority* is distributed: the labour inspectorate acts in employment contexts under occupational health and safety mandates, the national equality body addresses discrimination arising from refusal, and the DPA handles neural data processed during mandated use. No single existing authority possesses a comprehensive cognitive-liberty mandate; the IAP deliberately distributes enforcement across authorities with adjacent mandates.

The *remedy* comprises administrative orders prohibiting mandatory neurotechnology use without statutory authorisation, occupational health and safety enforcement for non-compliant workplace neural monitoring, discrimination remedies for individuals penalised for refusal, and DPA enforcement for neural data processed without valid legal basis. The IAP does not apply to voluntary consumer use of neurotechnology, clinical use prescribed by medical professionals with informed consent, or security-critical contexts where sector-specific legislation authorises cognitive monitoring (e.g., aviation safety) with statutory safeguards and independent oversight.

The IAP satisfies all five design principles: *verifiability* (contract review and procurement audit are institutionally standard); *institutional home* (distributed across labour inspectorate, equality body, and DPA); *trigger precision* (“conditioning access on neural device use” is observable from policy documents); *remedy adequacy* (enforcement uses existing administrative tools); *falsifiability* (voluntary use and clinical use are explicitly excluded). The normative invariant — freedom from *non-consensual* interference — is preserved through the focus on coercive contexts. The IAP narrows “cognitive liberty” from a comprehensive philosophical principle to an enforceable anti-coercion rule, accepting the semantic narrowing as the price of institutional traction.

## 6.5 Institutional Repair Paths and Trade-offs

Table 8: Institutional Repair Path Map: From Failure Mode to Governance Intervention

Failure Mode	Shortest Path (Soft Law)	Medium Path (Mandate Expansion)	Long Path (New Node)	Trade-offs
<b>VV-dominant</b> (high VV, low AF/TE): unverifiable	Develop technical measurement standards via ISO/IEC or CENELEC	Fund neurotechnology-specific auditing capacity within existing market surveillance authorities	Create dedicated neural-data auditing body	Standards may lag technology; auditing capacity is expensive
<b>AF-dominant</b> (low VV, high AF, low TE): measurable, no authority	Inter-authority MoU designating lead regulator for neurotechnology	Amend existing mandate of DPA or market surveillance authority to include neurotechnology explicitly	Create cross-cutting neurotechnology oversight agency	MoUs lack enforcement teeth; mandate amendment risks scope creep
<b>TE-dominant</b> (low VV, low AF, high TE): authority and evidence exist, no trigger	Issue sector-specific guidance specifying compliance thresholds	Codify thresholds in technical regulation (e.g., delegated acts under AI Act)	Enact neurotechnology-specific legislation with defined triggers	Premature threshold-setting may freeze technology; guidance lacks binding force
<b>Triple failure</b> (high VV+AF+TE): fully inert	Full downward translation via IAP design (this article’s method)	Staged translation: first secure institutional home (fix AF), then develop triggers (TE), then build verification infrastructure (VV)	Constitutional amendment with implementing statute and designated authority	Risk of premature operationalisation; constitutional concepts may lose generative power

*Precedent and time horizons.* VV-dominant repair has a precedent in the ISO/IEC JTC 1/SC 42 AI standards process (~3 years from NWIP to IS). AF-dominant repair is illustrated by the GDPR Article 70 EDPB mandate expansion (~2 years from proposal to codification). TE-dominant repair is expected

through the EU AI Act delegated acts process (2025–2027). Triple-failure repair remains an open challenge: Chile’s 2021 neurorights constitutional amendment has produced no implementing statute after >4 years, illustrating the difficulty of full-chain institutional construction.

The failure-mode typology closes the diagnosis-to-prescription loop: once the IAA identifies *where* a concept fails (VV, AF, TE, or all three), the repair path map identifies *what kind of institutional intervention* is required and at *what level of institutional ambition*. This makes the framework a genuine decision tool for governance designers, not merely a scoring exercise.

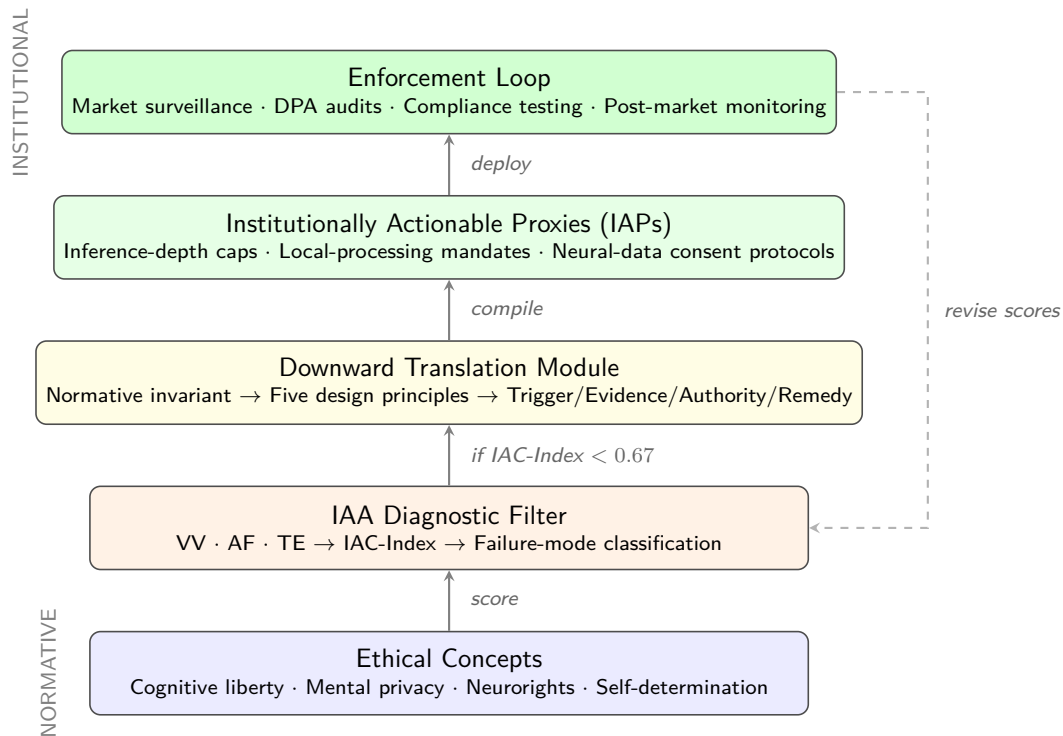


Figure 1: Ethics-to-Governance Translation Stack. Ethical concepts enter the IAA diagnostic filter, which scores their institutional readiness (VV/AF/TE). Concepts scoring below the actionability threshold undergo downward translation into Institutionally Actionable Proxies, which enter the enforcement loop. The feedback arrow indicates that enforcement experience revises IAA scores over time. This figure operationalises the structural logic linking conceptual ontology transformation to the operational requirements of institutional accountability.

## 7 Implications, Boundary Conditions, and Limitations

### 7.1 When Ethical Concepts Are Actionable (Negative Cases)

Not all ethical concepts require downward translation. This article’s framework applies to concepts that are institutionally inert *despite normative ambition*; it does not claim that all ethical concepts are inert. Four types of concepts achieve institutional actionability without the translation work described here:

**Type 1: Self-executing prohibition.** The prohibition of torture names a categorically identifiable act and applies universally without institutional discretion. No trigger threshold, evidentiary inference, or authority designation is required; the concept is self-executing. *Criterion:* the concept names an act that is categorically identifiable and universally prohibited.

**Type 2: Pre-operationalised right.** The GDPR right to erasure (Article 17) already specifies trigger (data subject request), authority (DPA), evidence (controller’s data records), and remedy (mandatory dele-

tion). The concept is fully operationalised in statutory text and requires no further translation. *Criterion:* the concept specifies trigger + authority + remedy in existing statutory text.

**Type 3: Constitutional aspiration / long-horizon norm.** “Dignity” and “autonomy” serve primarily as norm-formation and judicial-balancing concepts. Their institutional function is to provide interpretive anchors for courts and constitutional review bodies, not to drive enforcement action. Premature operationalisation would *reduce* their scope by locking them into specific triggers and thresholds that cannot anticipate future applications. *Criterion:* the concept’s primary function is norm-formation or judicial balancing, not enforcement.

**Type 4: Soft-law platform.** The OECD AI Principles (2019) are designed to enable future codification through iterative multi-stakeholder refinement. Locking such concepts into enforcement syntax too early forecloses legitimate design options and reduces the deliberative value of the soft-law process. *Criterion:* the concept is designed to enable future codification via iterative refinement.

**Application to the four neuroethics concepts:** Of the four concepts examined, *mental self-determination* most clearly belongs to Type 3 — it functions as a constitutional aspiration whose primary value lies in norm-formation and judicial interpretation, not enforcement. Premature operationalisation would risk trivialising a concept whose scope is legitimately open-ended. *Cognitive liberty* occupies a contested boundary between Type 3 and a concept that *should* be operationalised but *currently cannot be*: the liberal autonomy tradition demands enforceable protections, but the concept’s current formulation offers no institutional pathway. Mental privacy and neural data sovereignty are candidates for downward translation: they are not functioning as long-horizon norms, and their continued inertness produces real governance deficits.

## 7.2 Steelman Objections and Responses

### Objection 1: “Operationalising ethics cheapens it.”

This objection argues that reducing ethical concepts to measurable triggers and compliance thresholds strips them of their normative richness and moral force. The response is twofold. First, downward translation does not *replace* the ethical concept; it *supplements* it with an operational proxy. The philosophical concept of mental privacy continues to exist as a normative ideal; the IAP (inference-depth-capped local processing) exists as its institutional implementation. Second, a concept that remains permanently unimplemented is not “rich” — it is *inert*. Normative richness without institutional traction is indistinguishable from normative irrelevance. The five design principles (Section 6.2) and the normative invariant rule (Box 1) ensure that the translation preserves the concept’s irreducible semantic core.

### Objection 2: “Aspirational concepts serve a legitimate constitutional function.”

This objection is correct — and the boundary condition typology (Section 7.1) explicitly accommodates it. Type 3 concepts (constitutional aspirations) and Type 4 concepts (soft-law platforms) serve real governance functions through norm-formation, judicial interpretation, and deliberative refinement. The IAA framework does not claim that all concepts should be “compiled down.” It claims that concepts *which are deployed as if they provide governance protection but which lack institutional traction* create a governance mirage that is worse than honest admission of non-coverage. The framework provides the diagnostic tools to distinguish concepts that legitimately resist operationalisation (Type 3/4) from concepts that are merely unfinished institutional work.

### Objection 3: “The IAA privileges enforcement-friendly rights over justice-oriented rights.”

This objection poses the deepest challenge. If the IAA rewards concepts that are easy to enforce and penalises concepts that resist enforcement, it may systematically undervalue rights that protect against structural injustice, power asymmetry, or systemic discrimination — precisely the rights that are hardest to operationalise. The response is that the IAA is a *diagnostic* tool, not a *normative* one. A low IAC-Index

score does not mean a concept is unimportant; it means the concept is *not ready for institutional processing in its current form*. The appropriate response to a low score is not to abandon the concept but to undertake the translation work that would make it enforceable — while using the normative invariant check to ensure that the translation preserves what matters. The IAA identifies the work that needs to be done; it does not prejudge which concepts deserve that work.

### 7.3 The Risk of Premature Operationalisation

When should neuroethics concepts remain aspirational and resist premature operationalisation? Not every ethical concept should be “compiled down.” Three conditions favour leaving a concept institutionally inert:

*First*, when the concept’s primary function is *constitutional signalling*: establishing a normative floor that constrains future legislation without itself constituting enforcement (Type 3). Dignity serves this function; premature operationalisation of dignity into specific compliance tests would impoverish constitutional interpretation.

*Second*, when the technological landscape is evolving so rapidly that *any specific trigger or threshold would become obsolete within the regulatory revision cycle*. In such cases, soft-law platforms (Type 4) provide adaptive governance that rigid enforcement syntax cannot.

*Third*, when the concept’s *normative invariant cannot be preserved through operationalisation* — when any IAP that would satisfy VV, AF, and TE requirements would necessarily distort the concept’s irreducible core beyond recognition. This is a genuine risk for concepts like cognitive liberty, where the “cognitive” dimension resists the external verifiability that institutional processing requires.

### 7.4 Limitations and Research Agenda

The IAA framework as presented here has several limitations that define a research agenda. First, the scoring in this article is author-performed; future applications should employ multi-coder designs with formal inter-coder reliability assessment using the Box 1 protocol and the Cohen’s  $\kappa \geq 0.60$  threshold recommended therein. Second, the sensitivity analysis (Table 5) demonstrates rank-order invariance under three weighting schemes, but the space of possible weight configurations has not been exhaustively explored; future work should investigate whether non-linear or threshold-based aggregation rules improve diagnostic validity. Third, the cross-technology test (affective computing) is limited to a single concept; systematic application across multiple technology domains — digital therapeutics, synthetic biology, quantum computing governance — would strengthen the framework’s generalisability claim. Fourth, the two worked examples (mental privacy and cognitive liberty) demonstrate downward translation for a borderline and a triple-failure case respectively, but extending the exercise to neural data sovereignty and mental self-determination would complete the diagnostic picture for all four concepts. Fifth, the cross-jurisdiction comparison (Table 6) covers only two concepts across two jurisdictions; a full  $4 \times 3$  matrix (four concepts  $\times$  EU/US/Chile) would provide a more comprehensive test of architecture portability. Sixth, the framework evaluates institutional readiness at a single point in time (March 2026); longitudinal application tracking IAC-Index scores as regulatory frameworks evolve would test the framework’s sensitivity to institutional change.

## 8 Conclusion

Together, the translation paradox, the enforcement readability constraint, and the downward translation imperative specify the structural logic of emerging technology governance: ethical concepts are inherently unprocessable by regulatory syntax unless they are compiled into observable triggers, evidentiary

substrates, competent authorities, and enforceable remedies; governance therefore requires not more normative ambition but a translation module that converts aspiration into institutional format; and without explicit downward compilation — constrained by normative invariant checks to prevent semantic loss — concept proliferation structurally forecloses effective regulation by creating the appearance of coverage where none exists. The IAA framework provides that translation module, the IAC-Index measures the distance still to travel, and the five design principles constrain the engineering so that what institutions can process remains faithful to what ethics demands.

### **Two theoretical insights.**

*First*, the ethics-governance relationship is not one of progressive realisation — where ethical concepts gradually mature into enforceable rules — but one of *structural discontinuity*. Ethical concepts and institutional governance operate in different grammars with different success conditions. The same properties that make a concept normatively compelling (abstraction, comprehensiveness, philosophical depth) are precisely the properties that make it institutionally unprocessable. The “translation” from ethics to governance is therefore not a matter of patience or political will; it requires active engineering of operational proxies that preserve normative intent in a format institutions can consume.

*Second*, institutional actionability is a *relational* property, not an intrinsic one. A concept is not actionable or inert *per se*; it is actionable *for a specific institution in a specific jurisdiction under specific regulatory conditions*. The IAA framework’s insistence on the concept  $\times$  institution  $\times$  jurisdiction triple prevents the reification of actionability as a fixed property and forces analysts to specify the institutional context in which their assessments hold. This relational ontology has implications beyond neurotechnology: it suggests that debates about whether ethical concepts “work” in governance are systematically under-specified unless they name the institutional-jurisdictional context.

### **Three research questions.**

*First*, can the IAA be validated through multi-coder empirical application? The framework’s replicability claim requires testing: multiple research teams scoring the same concepts in the same institutional contexts should produce convergent IAC-Index scores. A multi-site coding study using the IAA codebook and the Box 1 protocol would constitute a strong validity test.

*Second*, how do IAC-Index scores change as regulatory frameworks evolve? A longitudinal study tracking IAC-Index scores for the four neuroethics concepts across the EU, US, and Chilean institutional contexts over a five-year period (2026–2031) would test whether legislative and regulatory developments (the EU AI Act, Chilean implementing legislation, ICO neurodata guidance) produce measurable changes in institutional actionability. Significant score changes would validate the framework’s sensitivity; stable scores despite regulatory activity would suggest that the actionability gap is more deeply structural than legislative.

*Third*, does the IAA framework transfer effectively to non-neurotech domains? Systematic application to affective computing (emotional privacy, algorithmic emotional manipulation), synthetic biology (biosafety sovereignty, genetic self-determination), and quantum computing governance (cryptographic obsolescence, quantum information rights) would test whether the VV/AF/TE indicators capture institutional readiness dynamics across technology domains or are implicitly calibrated to neurotechnology-specific governance structures.

### **One trade-off.**

Concepts that cannot be parsed into triggers, evidence, authorities, and remedies are governance aspirations, not governance instruments. But governance aspirations are not worthless — they establish normative horizons, constrain future legislation, and provide interpretive resources for courts and constitutional bodies. The trade-off inherent in downward translation is the loss of this generative ambiguity: a concept that has been compiled into an IAP gains enforcement teeth but loses the open-ended normative scope that allows it to evolve with changing social conditions and technological possibilities. The gover-

nance challenge for neurotechnology — and for every emerging technology that outpaces institutional capacity — is to navigate this trade-off deliberately rather than by default: to identify which concepts should be compiled down, which should remain aspirational, and to make that choice with full awareness of what is gained and what is lost.

## References

- Baldwin, R., Cave, M., & Lodge, M. (2012). *Understanding regulation: Theory, strategy, and practice* (2nd ed.). Oxford University Press.
- Biasiucci, A., Franceschiello, B., & Murray, M. M. (2019). Electroencephalography. *Current Biology*, 29(3), R80–R85.
- Black, J. (2002). Critical reflections on regulation. *Australian Journal of Legal Philosophy*, 27, 1–35.
- Black, J. (2008). Constructing and contesting legitimacy and accountability in polycentric regulatory regimes. *Regulation & Governance*, 2(2), 137–164.
- Bradford, A. (2020). *The Brussels effect: How the European Union rules the world*. Oxford University Press.
- Bublitz, J. C. (2013). My mind is mine!? Cognitive liberty as a legal concept. In E. Hildt & A. G. Franke (Eds.), *Cognitive enhancement: An interdisciplinary perspective* (pp. 233–264). Springer.
- Bublitz, J. C. (2015). Cognitive liberty or the international human right to freedom of thought. In J. Clausen & N. Levy (Eds.), *Handbook of neuroethics* (pp. 1309–1333). Springer.
- Burwell, S., Sample, M., & Racine, E. (2017). Ethical aspects of brain computer interfaces: A scoping review. *BMC Medical Ethics*, 18(1), 60.
- Büthe, T., & Mattli, W. (2011). *The new global rulers: The privatization of regulation in the world economy*. Princeton University Press.
- Bygrave, L. A. (2014). *Data privacy law: An international perspective*. Oxford University Press.
- Callon, M. (1986). Some elements of a sociology of translation: Domestication of the scallops and the fishermen of St Briec Bay. In J. Law (Ed.), *Power, action and belief: A new sociology of knowledge?* (pp. 196–233). Routledge & Kegan Paul.
- Cohen, J. E. (2013). What privacy is for. *Harvard Law Review*, 126(7), 1904–1933.
- Colorado General Assembly. (2024). House Bill 24–1058: Protections for biological and neural data under the Colorado Privacy Act. Effective August 7, 2024.
- Council of Europe. (2025). Draft guidelines on data protection in the neurosciences. Consultative Committee of Convention 108+, September 2025.
- Dworkin, R. (1977). *Taking rights seriously*. Harvard University Press.
- European Data Protection Supervisor. (2024). TechDispatch #1/2024 — Neurodata. EDPS, 3 June 2024.
- European Data Protection Supervisor & Agencia Española de Protección de Datos. (2024). Joint report on the data protection challenges of neurotechnology. EDPS & AEPD, 27 June 2024.
- European Parliament & Council. (2016). Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation). *Official Journal of the European Union*, L 119/1.
- European Parliament & Council. (2017). Regulation (EU) 2017/745 on medical devices (Medical Devices Regulation). *Official Journal of the European Union*, L 117/1.
- European Parliament & Council. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L 2024/1689.
- Farahany, N. A. (2023). *The battle for your brain: Defending the right to think freely in the age of neurotechnology*. St. Martin's Press.
- U.S. Food and Drug Administration. (2019). *General wellness: Policy for low risk devices — Guidance for industry and Food and Drug Administration staff* (updated). FDA.

- Fins, J. J. (2015). *Rights come to mind: Brain injury, ethics, and the struggle for consciousness*. Cambridge University Press.
- Fitz, N. S., & Reiner, P. B. (2015). The challenge of crafting policy for do-it-yourself brain stimulation. *Journal of Medical Ethics*, 41(5), 410–412.
- Fleischer, V. (2010). Regulatory arbitrage. *Texas Law Review*, 89(2), 227–289.
- Floridi, L. (2018). Soft ethics and the governance of the digital. *Philosophy & Technology*, 31(1), 1–8.
- Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
- Ford, C. L. (2017). *Innovation and the state: Finance, regulation, and justice*. Cambridge University Press.
- Frank, L. E., Brinke, J. K., & Schermer, M. H. N. (2022). Direct-to-consumer neurotechnology: What is it and what is it for? *Neuroethics*, 15(1), 11.
- Garden, H., Winickoff, D. E., Frahm, N. M., & Pfothauer, S. (2019). Responsible innovation in neurotechnology enterprises. *OECD Science, Technology and Industry Working Papers*, 2019/05.
- Gilad, S. (2010). It runs in the family: Meta-regulation and its siblings. *Regulation & Governance*, 4(4), 485–506.
- Goering, S., & Yuste, R. (2016). On the necessity of ethical guidelines for novel neurotechnologies. *Cell*, 167(4), 882–885.
- Hood, C., Rothstein, H., & Baldwin, R. (2001). *The government of risk: Understanding risk regulation regimes*. Oxford University Press.
- Hoofnagle, C. J. (2016). *Federal Trade Commission privacy law and policy*. Cambridge University Press.
- Information Commissioner's Office. (2023). ICO tech futures: Neurotechnology. ICO, June 2023.
- Information Commissioner's Office. (2026). Technology guidance roadmap: Neurotechnology and neurodata guidance. ICO.
- IEEE. (2024). *IEEE Neuroethics Framework*. IEEE Standards Association.
- Ienca, M., & Andorno, R. (2017). Towards new human rights in the age of neuroscience and neurotechnology. *Life Sciences, Society and Policy*, 13(1), 5.
- Ienca, M., & Vayena, E. (2020). On the responsible use of digital data to tackle the COVID-19 pandemic. *Nature Medicine*, 26(4), 463–464.
- Ienca, M. (2021). On neurorights. *Frontiers in Human Neuroscience*, 15, 701258.
- Illes, J., & Raffin, T. A. (2006). No child left without a brain scan? Toward a pediatric neuroethics. *Cerebrum*, 2006, 1–14.
- Jasanoff, S. (2004). The idiom of co-production. In S. Jasanoff (Ed.), *States of knowledge: The co-production of science and the social order* (pp. 1–12). Routledge.
- Jasanoff, S. (2016). *The ethics of invention: Technology and the human future*. W. W. Norton.
- Jwa, A. S., & Poldrack, R. A. (2022). The spectrum of data sharing policies in neuroimaging data repositories. *Human Brain Mapping*, 43(8), 2707–2721.
- Kellmeyer, P. (2021). Big brain data: On the responsible use of brain data from clinical and consumer-directed neurotechnological devices. *Neuroethics*, 14, 83–98.
- Kellmeyer, P., Cochrane, T., Müller, O., Mitchell, C., Ball, T., Fins, J. J., & Biller-Andorno, N. (2016). The effects of closed-loop medical devices on the autonomy and accountability of persons and systems. *Cambridge Quarterly of Healthcare Ethics*, 25(4), 623–633.
- Klein, E., & Ojemann, J. (2016). Informed consent in implantable BCI research. *Journal of Neural Engineering*, 13(4), 043001.
- Levi-Faur, D. (2011). Regulation and regulatory governance. In D. Levi-Faur (Ed.), *Handbook on the politics of regulation* (pp. 3–21). Edward Elgar.
- Levi-Faur, D. (2012). From “big government” to “big governance”? In D. Levi-Faur (Ed.), *The Oxford handbook of governance* (pp. 3–18). Oxford University Press.

- Lodge, M., & Wegrich, K. (2012). *Managing regulation: Regulatory analysis, politics, and policy*. Palgrave Macmillan.
- Lynskey, O. (2015). *The foundations of EU data protection law*. Oxford University Press.
- Majone, G. (1997). From the positive to the regulatory state: Causes and consequences of changes in the mode of governance. *Journal of Public Policy*, 17(2), 139–167.
- Marchetti, B. (2021). Digital governance meets regulatory governance: A framework for analysis. *European Journal of Risk Regulation*, 12(2), 291–303.
- Martinovic, I., Davies, D., Frank, M., Perito, D., Ros, T., & Song, D. (2012). On the feasibility of side-channel attacks with brain–computer interfaces. In *Proceedings of the 21st USENIX Security Symposium* (pp. 143–158). USENIX Association.
- Maslen, H., Douglas, T., Cohen Kadosh, R., Levy, N., & Savulescu, J. (2014). The regulation of cognitive enhancement devices: Extending the medical model. *Journal of Law and the Biosciences*, 1(1), 68–93.
- Minnesota Legislature. (2025). Senate File 1240: Neurodata rights and mental privacy. 2025–2026 session.
- Moran, M. (2002). Understanding the regulatory state. *British Journal of Political Science*, 32(2), 391–413.
- Müller, O., & Rotter, S. (2017). Neurotechnology: Current developments and ethical issues. *Frontiers in Systems Neuroscience*, 11, 93.
- Nussbaum, M. C. (2011). *Creating capabilities: The human development approach*. Harvard University Press.
- OECD. (2019). *Recommendation of the Council on Responsible Innovation in Neurotechnology*. OECD/LEGAL/0457.
- Ostrom, E. (2010). Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100(3), 641–672.
- Prosser, T. (2010). *The regulatory enterprise: Government, regulation, and legitimacy*. Oxford University Press.
- Richards, N. M. (2013). The dangers of surveillance. *Harvard Law Review*, 126(7), 1934–1965.
- Rommelfanger, K. S., Jeong, S. J., Ema, A., Fukushi, T., Kasai, K., Ramos, K. M., ... & the Global Neuroethics Summit Delegates. (2018). Neuroethics questions to guide ethical research in the international brain initiatives. *Neuron*, 100(1), 19–36.
- Roskies, A. (2002). Neuroethics for the new millennium. *Neuron*, 35(1), 21–23.
- Salles, A., Bjaalie, J. G., Evers, K., Farisco, M., Fothergill, B. T., Guerrero, M., ... & Bhatt, S. (2019). The Human Brain Project: Responsible brain research for the benefit of society. *Neuron*, 101(3), 380–384.
- Schwartz, P. M., & Solove, D. J. (2011). The PII problem: Privacy and a new concept of personally identifiable information. *New York University Law Review*, 86(6), 1814–1894.
- Scott, C. (2004). Regulation in the age of governance: The rise of the post-regulatory state. In J. Jordana & D. Levi-Faur (Eds.), *The politics of regulation* (pp. 145–174). Edward Elgar.
- Sen, A. (1999). *Development as freedom*. Alfred A. Knopf.
- Senate of Chile. (2021). Ley 21.383: Modifica la Carta Fundamental, para establecer la protección de los neuroderechos o derechos del cerebro humano. Diario Oficial de la República de Chile.
- Shen, F. X. (2013). Neuroscience, mental privacy, and the law. *Harvard Journal of Law and Public Policy*, 36(2), 653–713.
- Solove, D. J. (2013). Introduction: Privacy self-management and the consent dilemma. *Harvard Law Review*, 126(7), 1880–1903.
- Sparrow, M. K. (2000). *The regulatory craft: Controlling risks, solving problems, and managing compliance*. Brookings Institution Press.
- Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, “translations” and boundary objects: Amateurs and professionals in Berkeley’s Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science*, 19(3), 387–420.
- Stilgoe, J., Owen, R., & Macnaghten, P. (2013). Developing a framework for responsible innovation. *Research Policy*, 42(9), 1568–1580.

- Tene, O., & Polonetsky, J. (2012). Big data for all: Privacy and user control in the age of analytics. *Northwestern Journal of Technology and Intellectual Property*, 11(5), 239–272.
- UNESCO. (2023). *Recommendation on the ethics of neurotechnology*. UNESCO, November 2023.
- Wexler, A. (2017). The social context of “do-it-yourself” brain stimulation: Neurohackers, biohackers, and lifehackers. *Frontiers in Human Neuroscience*, 11, 224.
- Wexler, A., & Thibault, R. (2019). Mind-reading or misleading? Assessing direct-to-consumer electroencephalography (EEG) devices marketed for wellness and their ethical and regulatory implications. *Journal of Cognitive Enhancement*, 3(1), 131–137.
- Winner, L. (1980). Do artifacts have politics? *Daedalus*, 109(1), 121–136.
- Yuste, R., Goering, S., Arcas, B. A. Y., Bi, G., Carmena, J. M., Carter, A., ... & Wolpaw, J. (2017). Four ethical priorities for neurotechnologies and AI. *Nature*, 551(7679), 159–163.
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.