ie UNIVERSITY
CENTER FOR THE GOVERNANCE OF CHANGE

PAPER #3
OCTOBER 2026



cgc.ie.edu

THE NEUROTECHNOLOGY SHIFT

The Madrid Scale for Neurotechnology Governance

A Use-Based Calibration of Risk and Opportunity in Context

This paper was prepared for the Center for the Governance of Change (CGC) at IE University as part of The Neurotechnology Shift, a research program of the CGC and Fundación Tatiana. It draws on three sources:

(i) academic and policy literature on neurotechnology governance, neurodata, and ecological approaches to cognition;

(ii) institutional and regulatory materials, including the OECD Recommendation on Responsible Innovation in Neurotechnology, the UNESCO Recommendation on the Ethics of Neurotechnology, and the EU AI Act, General Data Protection Regulation, and Medical Device Regulation; and

(iii) a closed-door expert consultation convened by the CGC over a working lunch on September 30, 2026, which brought together twenty participants from government, regulatory authorities, industry, and academia[1]. The discussion was conducted under the Chatham House Rule, and where this paper reflects perspectives expressed there, it does so without attribution to individual participants. The analysis and recommendations are the author's own.

This paper presents the ***Madrid Scale for Neurotechnology Governance v1.0 (October 2026)***, the first version of a framework designed to evolve through further multi-stakeholder validation. The Scale is illustrative and directional, not prescriptive or legally binding: its profiles and levels are meant to structure judgment, not to replace the decisions of competent authorities. Each version carries a number and a release date. A new first digit (v2.0) signals a change to the dimensions, the calibration levels, or the reading rule, so assessments made under different versions are not directly comparable; a new second digit (v1.1) refines guidance without changing how profiles are read.

AUTHOR:
José M. Muñoz

TABLE OF CONTENTS

KEY ELEMENTS 4

1. INTRODUCTION: FROM PRINCIPLES TO IMPLEMENTATION 6

2. WHY EXISTING INSTRUMENTS DO NOT DECIDE CASES 7

3. THE MADRID SCALE: AN ECOLOGICAL APPROACH TO GOVERNING USE 8

4. READING THE SCALE 10
   - 4.1. Calibration levels 11
   - 4.2. From dimensions to calibrated use 12
   - 4.3. Recalibration and review 14

5. ILLUSTRATING THE SCALE: ONE HEADBAND, THREE USES 15

6. APPLYING THE MADRID SCALE: FIVE STEPS AND RECOMMENDATIONS FOR DECISION-MAKERS 17
   - 6.1. Applying the Scale in five steps 17
   - 6.2. Recommendations for decision-makers 18

7. NEXT STEPS: BUILDING ON THIS VERSION 19

8. FINAL REMARKS: GOVERN THE USE, NOT THE SOURCE 20

9. THE MADRID SCALE REFERENCE CARD 21

REFERENCES AND NOTES 23



KEY ELEMENTS

The challenge: From principles to implementation

Neurotechnology has moved beyond the clinic into homes, workplaces, and classrooms, and AI is widening what brain data can reveal once they are combined with other personal data. The OECD and UNESCO have laid a strong foundation, setting out what deserves protection and the actions available to policymakers, while binding rules such as the EU AI Act, the GDPR, and the MDR classify devices, data categories, or declared purposes. The next step is to decide, use by use, whether a specific use can proceed under existing rules or calls for more. That decision needs a threshold.

The proposal: Govern the use, not the source

The Madrid Scale is an ecological, use-based framework for calibrating neurotechnology governance to risk and opportunity in context. Its unit of analysis is the use: the same headband required on a flight deck, worn by choice at home, or scored in a classroom receives three different profiles. The Scale does not compete with existing law but maps onto it, showing where current rules suffice and where more is needed. Its outputs are illustrative and directional, not prescriptive or legally binding.

How it works: Four dimensions, four levels, one reading rule

Each use is examined along four ecological dimensions:



Embodied (what the signal reveals)



Embedded (where, by whom, on whom, and for what purpose)



Enactive (what changes in the person and the system with use)



Extended (what the use connects to and what returns to the person)

Each dimension weighs risk and opportunity and reaches one of four calibration levels: *Standard, Safeguarded, Controlled,* or *Untenable use*. The highest dimension sets the label, never an average, so a critical concern cannot be offset elsewhere. Consent, accountability, and reversibility are cross-cutting conditions that can raise a dimension or make it harder to lower, but never lower it. Every calibration is reviewed when evidence, context, or deployment scale changes.

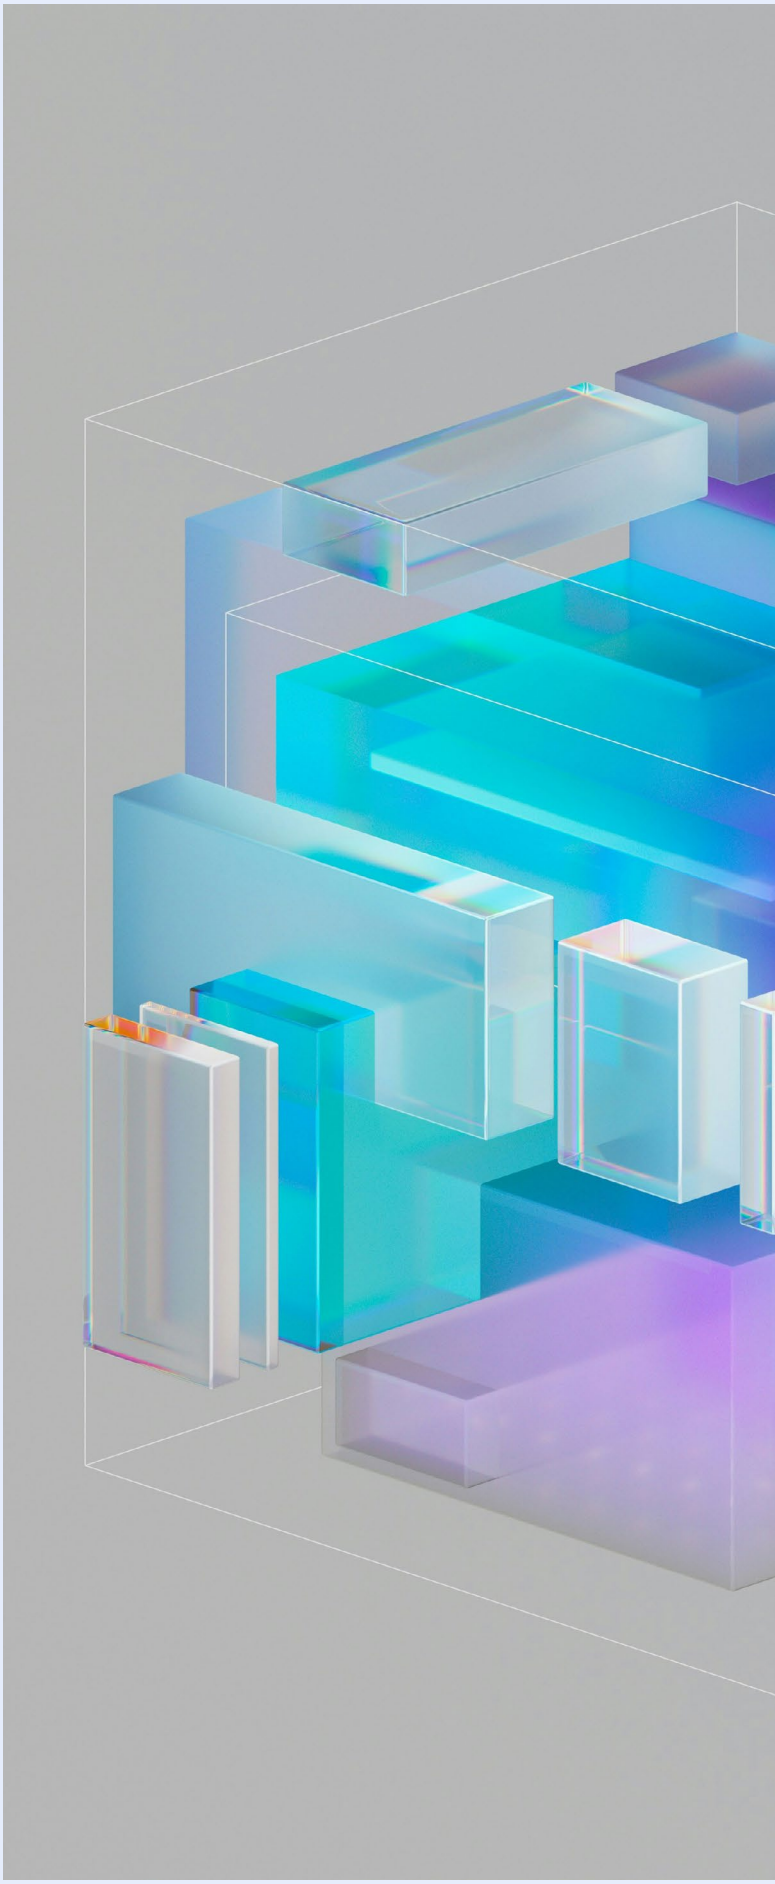
Who applies it and how

The Scale applies to every use, clinical and research ones included, and fits into decisions already made at design, funding, approval, procurement, and after deployment. Five steps lead from defining the use to a recorded assessment with its review regime. Existing assessments count as evidence instead of being repeated, and a repurposed device needs a new assessment by whoever deploys the new use. Because profiles describe uses rather than laws, authorities in different jurisdictions can read them in common terms. The paper closes with recommendations for regulators, public funders, ethics committees and procurement, industry, and international organizations.

The Madrid Scale Reference Card: A ready-to-go tool for decision-makers

A two-page Reference Card accompanies this paper for those who apply the Scale in practice. It gathers what an assessment needs at the point of decision: the four dimensions and the questions they ask, the calibration levels with their governance responses, and the floors set by three cross-cutting conditions. It also shows how to read a profile and sets out the five steps for applying the Scale. Twelve questions, which test each cross-cutting condition in each dimension, serve as a working checklist.

The Reference Card is available ▸ here.



1. INTRODUCTION: FROM PRINCIPLES TO IMPLEMENTATION

Neurotechnology has left the clinic; the **neurotechnology shift** is here[2]. Devices that record or modulate brain activity now reach homes, offices, and classrooms, often as consumer products subject to lighter rules than their medical predecessors, as the earlier papers in *The Neurotechnology Shift* project series have shown[3]. Artificial intelligence is widening what these devices can reveal, especially once brain signals are merged with other personal data[4]. Even though not all neurotechnology depends on AI, the convergence helps explain why a single device can carry very different consequences depending on where and why it is used. This calls less for alarm than for precision.

Policy has moved quickly. UNESCO's Member States adopted the first global Recommendation on the Ethics of Neurotechnology in November 2025[5], and the MIND Act, introduced in the US Senate that year, would ask the Federal Trade Commission to examine how neurodata is used[6]. The European Brain Council is piloting a voluntary Seal of Responsible Neurotechnology[7], and both the 1st CINET Neurotechnology Conference[8] and the 2026 Asilomar for the Brain and Mind summit[9] sought working tools rather than another charter.

These efforts unfold amid strategic competition[10], and jurisdictions increasingly part ways on how they classify devices and data. A reference built around uses, not around any one legal system, can help them read one another. Importantly, the scarce resource is not principles. Mental privacy, integrity, and autonomy are widely endorsed, and the field's task is shifting toward making such commitments operational across very different settings[11]. Yet a principle cannot tell a funder or a regulator whether a particular use may proceed under existing rules or calls for something more, because that decision needs a threshold.

This paper takes that step **from principles to implementation**. The Madrid Scale is an ecological, use-based framework for calibrating neurotechnology governance to risk and opportunity in context. Its unit of analysis is the use, not the device: the same headband worn by choice at home, required on a flight deck, or scored in a classroom receives a different profile. For each use, the Scale shows where attention is needed, which response fits, and how the use can move as its design or the evidence changes. Its outputs are illustrative and directional, not prescriptive or legally binding.



2. WHY EXISTING INSTRUMENTS DO NOT DECIDE CASES

International tools such as the OECD 2019 Recommendation on Responsible Innovation in Neurotechnology[12] and UNESCO's 2025 Recommendation on the Ethics of Neurotechnology[13] set out what deserves protection, and the OECD's 2025 Neurotechnology Toolkit adds a range of possible actions[14]. These instruments established guidelines for responsible innovation in neurotechnology, laying a very important foundation. Building on this work, governance of specific uses is a natural next step towards a more comprehensive framework. Where the Toolkit sets out possible actions, the **Madrid Scale for Neurotechnology Governance** helps identify which response fits a given use in a given context.

Binding regulation, for its part, usually classifies by device type, by category of data, or by the purpose a manufacturer declares. Under the EU Medical Device Regulation (MDR)[15], a recording headband is covered only if marketed for a medical purpose, whereas transcranial electrical or magnetic stimulation equipment is covered even without one. The General Data Protection Regulation (GDPR) protects health and biometric data but does not name neurodata[16]. Many devices therefore sit in a regulatory grey zone at the boundary of wellness and medicine[17]. The difficulty lies less in the amount of regulation than in where it is aimed.

Regulation is not alone in this. The rapid evolution of neurotechnology and the expanding possibilities it offers are creating new governance challenges. The maps the field uses to describe itself sort technologies by modality and by whether they are invasive or non-invasive[18], and none of them says whether a given use, in a given setting, is acceptable. Even their main axis is losing its edges: telling invasive from non-invasive methods is getting harder[19], since techniques that never enter the skull can still act on the brain or collect highly sensitive data, and devices keep crossing the categories designed to contain them, from the clinic to the consumer market[20].

Chile, which amended its constitution in 2021 to protect brain activity and the information derived from it[21], shows the same gap from the side of principles. In 2023, ruling on a consumer electroencephalography (EEG) headset, its Supreme Court cited that protection but decided the case on existing privacy and data protection grounds and referred the device to the health and customs authorities[22]. The principle was in place; the threshold to apply it was still not there.

A potential objection is that the EU AI Act already accounts for context: it prohibits emotion recognition in workplaces and educational institutions and designates certain uses as high-risk[23]. Its context, however, is a predefined list that applies only to AI systems and starts from the purpose the provider declares. It does not reach uses without AI, imposes few obligations on those it treats as minimal risk, and engages with a change of setting only when the new use falls into one of its listed categories.

The Madrid Scale **does not compete with the AI Act, the GDPR, or the MDR** but maps onto them, indicating when existing rules suffice and when more is needed. Europe is one example among several, and the same reasoning applies to US state privacy laws[24] or to Chile's constitutional route.

3. THE MADRID SCALE: AN ECOLOGICAL APPROACH TO GOVERNING USE

The Madrid Scale assesses uses, not devices. A use is an application in a given context and within the data ecosystem it connects to. The same headband is a different use at home, at work, or in a classroom, and each context can be divided further as a decision requires.

Each use is examined along **four ecological dimensions**—grounded in how a neurotechnology actually works in everyday life, rather than in the laboratory or in what its maker declares. Each one starts from a premise about how neurotechnology operates and turns it into a plain question that weighs risk and opportunity together. Their names come from 4E-cognition theories (embodied, embedded, enactive, and extended), which hold that the mind depends on the body, the setting, ongoing activity, and external tools, not only on the brain[25]:



Embodied neurotechnology governance understands neural signals as part of the body's wider biometric ecosystem, alongside heart rate, eye movements, or voice, all of which can support mental inferences[26]. It asks what the signal reveals by means of neurodata inference, alone or combined with other bodily data. The more a signal reveals, the greater both its value for health and its potential for misuse, depending on what can actually be inferred and how easily people can be re-identified[27].



Embedded neurotechnology governance understands that the same device carries very different risks and opportunities depending on where it is used and for whom. It asks where the use takes place, who deploys it, on whom, and for what purpose. Here the Madrid Scale considers vulnerability: children, patients, or workers with little bargaining power face the same device on very different terms[28]. The device is constant; the stakes are not.



Enactive neurotechnology governance understands that person and system shape each other through use, so that risks and benefits keep emerging after a product reaches the market. It asks what changes, in the person and in the system, as the use continues. This is clearest in adaptive and closed-loop systems[29], but it also happens when people change their habits, self-understanding, or decisions in response to any neurotechnological device. Effects therefore depend on the history of use and cannot be fully characterized before deployment. Pre-market review, where it exists, is necessary but not sufficient, and governance must also follow the use.



Extended neurotechnology governance understands the brain as one node in an ecosystem of devices, data, and AI systems in which information flows both ways[30]: neurodata feed digital systems, and those systems act back on the person through recommendations, feedback, or stimulation. It asks what the use connects to, where its data and inferences go, and what returns to the person. Neurotechnology increasingly operates within AI-enabled infrastructures[31] in which neurodata are aggregated with behavioral and digital data, used to train models, and fed back to shape what the person sees or experiences. This coupling creates risks of profiling and of reuse beyond the original purpose, but data interoperability and responsible sharing[32] also create opportunity, from early biomarkers to better treatment.

Together, the four dimensions describe a single ecological approach to neurotechnology governance, not four separate compartments. Approaches to embodied, embedded, enactive, and extended cognition overlap and share ecological principles[33], even as they disagree on important points. The Madrid Scale separates them for analysis, not because they act independently. An AI system that combines brain signals with other data (Extended) changes what those signals reveal (Embodied), what a given setting can do with them (Embedded), and how the person and the system change as the use continues (Enactive). Because a problem in one dimension affects the others, a high level in one dimension cannot be offset by low levels elsewhere. That is the basis for the reading rule in Section 4.

Additionally, **three cross-cutting conditions** are examined within every dimension:

- **Consent** asks whether the user knows what is happening and can genuinely refuse. Because a use keeps changing after deployment, consent cannot be a single event. Dynamic consent[34], which lets people review and withdraw permission as their data find new purposes, suits neurotechnology well, although consent that can be withdrawn is still not free where refusing carries an unbearable cost.
- **Accountability** asks whether someone identifiable answers for the use and offers a route for complaint.
- **Reversibility** asks whether potential harm can be undone, both in the person and in the data, including data already used to train models.

These conditions can raise a dimension, or make it harder to lower, but never lower it, as Section 4 explains. **Table 1** turns them into twelve questions, one for each condition in each dimension.

Table 1.
Twelve questions: The cross-cutting conditions for neurotechnology uses in each governance dimension

	Consent	Accountability	Reversibility
EMBODIED	Does the person know what is recorded and what can be inferred from it?	Who answers for the accuracy of what is inferred?	Can recorded data and inferences be deleted?
EMBEDDED	Can the person refuse without losing something they value, such as a job, a grade, or a service?	Who deploys the use, and where can those affected complain?	If the use feeds a decision about the person (e.g., a job, a grade, a service), can that decision be contested and reversed?
ENACTIVE	Is the person told how the system adapts and what it may change in them?	Who monitors effects over time and acts on them?	Do changes in the person fade when the use stops?
EXTENDED	Does the person know where the data go and agree to each further use?	Who answers for third parties that receive data or inferences?	Can shared data be recalled, including from trained models?

4. READING THE SCALE

The profile places a use on the four dimensions of ecological neurotechnology governance:

- Embodied *(top)*
- Embedded *(left)*
- Enactive *(right)*
- Extended *(bottom)*

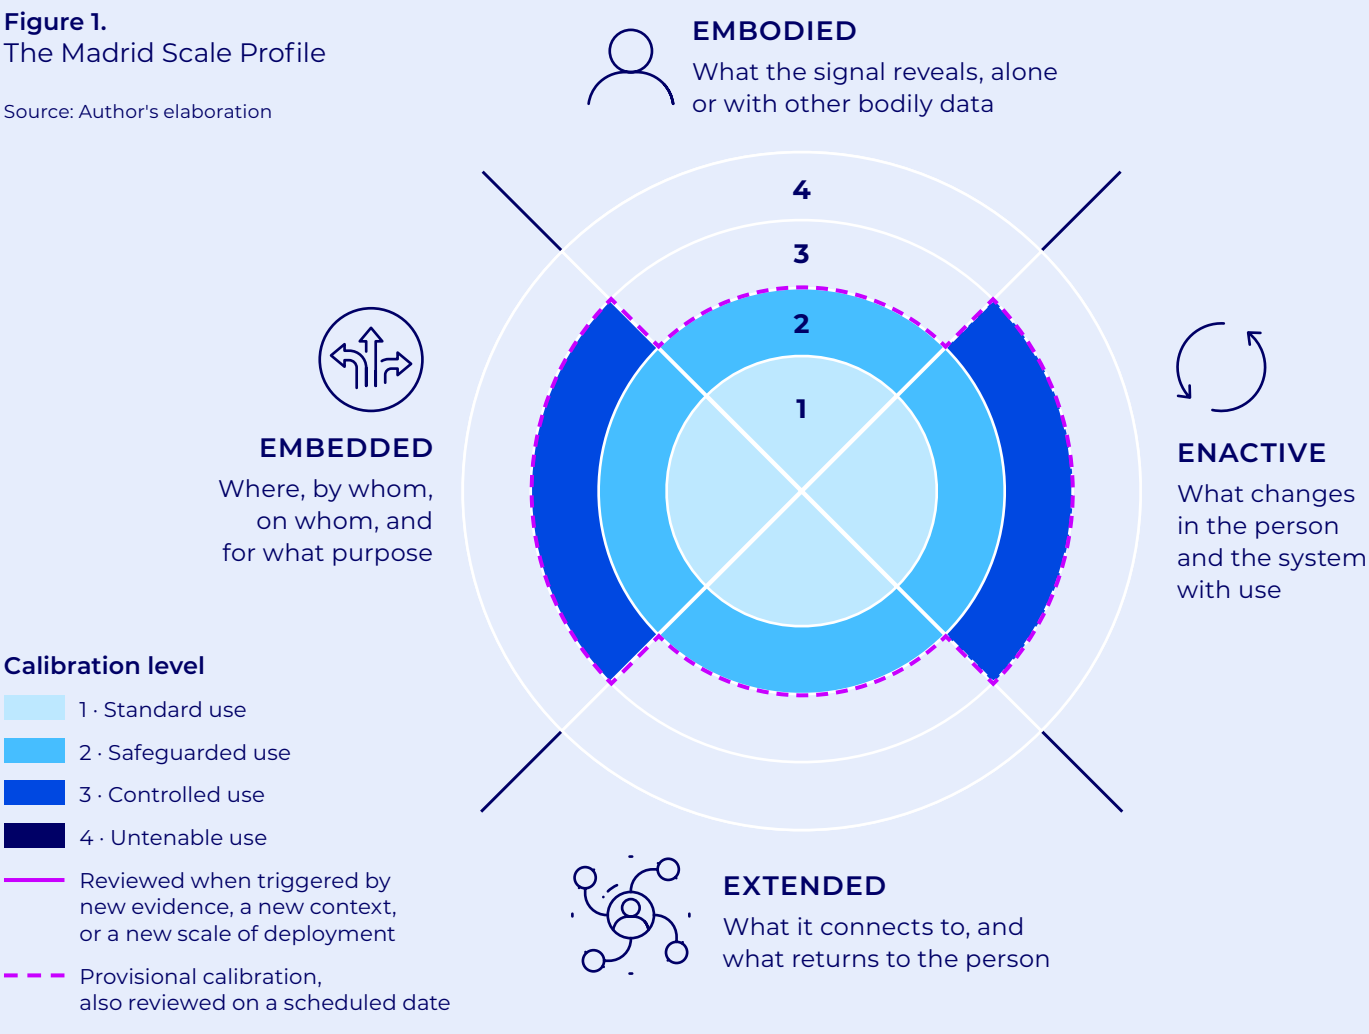
the order in which a profile is written.

Each level a dimension reaches adds a block, shaded from light (level 1) to dark (level 4), so the dimension that sets the label carries the darkest block.

The profile shown does not assess any real product, and is illustrative: *Controlled use* (2-3-3-2), in which the Embedded and Enactive dimensions share the highest level and together set the label.

Figure 1.
The Madrid Scale Profile

Source: Author's elaboration



CALIBRATION LEVEL	GOVERNANCE RESPONSE
1 · Standard use	Existing frameworks apply
2 · Safeguarded use	Added safeguards (e.g., standards, certification, limits on data use)
3 · Controlled use	Controlled deployment (e.g., pilot, sandbox, monitoring) while evidence is generated
4 · Untenable use	Not recommended in this form; only a change of use can move it

Reading the profile

1. Each dimension weighs risk and opportunity within itself.
2. The highest dimension sets the label, never an average or a sum.
3. *Label* + profile in fixed order (Embodied-Embedded-Enactive-Extended). Example: *Controlled use* (2-3-3-2).
4. Consent and accountability can raise a dimension; reversibility can make it harder to lower. None can lower it.
5. Illustrative and directional, not prescriptive or legally binding.

Review regime

— **Solid outline:** reviewed when triggered by new evidence, a new context, or a new scale of deployment.

- - - **Dashed outline (as in this example):** provisional calibration, also reviewed on a scheduled date.

4.1 CALIBRATION LEVELS

The Madrid Scale Profile **(Figure 1)** sets out the **four calibration levels**, numbered 1 to 4, with the governance response each calls for.

01 **Level 1** *(Standard use)* relies on existing frameworks, which already cover its risks.

02 **Level 2** *(Safeguarded use)* adds specific measures, such as standards, certification, or limits on data use, for risks that general rules do not target.

03 **Level 3** *(Controlled use)* applies when evidence suffices to test a use but not to scale it, so a pilot, sandbox, or monitored deployment pursues the benefit while generating that evidence.

04 **Level 4** *(Untenable use)* differs in kind, because the problem lies in the use itself rather than in missing evidence. This is the case when people are profiled without knowing it, are exposed to a use whose purpose is illegitimate, or cannot refuse without a cost they cannot bear. No safeguard repairs such a use and no controlled study makes it acceptable, since knowing how well it works does not change what it does to the people subject to it.

Human rights set the limit here[35]: the use is not recommended in this form, review does not soften that conclusion, and only a change of use could move it. Level 4 is therefore the only point at which the Scale advises against a use instead of calibrating how it proceeds.

Because some failures are inevitable, measures at every level should also limit harm, through reversibility, data deletion, or routes for complaint. A recent proposal for the governance of brain data calibrates protections to risk in a similar way[36].

4.2 FROM DIMENSIONS TO CALIBRATED USE

Each dimension is placed at a level by **weighing risk against** opportunity **(Table 2)**, and opportunity counts only under three rules. First, it must be *demonstrated*, because a promised benefit cannot offset a known risk. Second, it is *measured against the best alternative*, including those that collect no neural data, because exposure is unjustified if the same benefit is available without it. Third, it depends on *who gains and who bears the risk*: a benefit to whoever imposes a risk weighs less than protection for those who carry it[37].

The **cross-cutting conditions** then set floors, because each level presupposes something:

- **Consent.** Partial consent keeps the affected dimension at level 2 *(Safeguarded use)* or above, since *Standard use* assumes informed and freely given consent. If consent fails altogether, through covert inference or a cost of refusal the person cannot bear, nothing can repair it, and the dimension reaches level 4 *(Untenable use)*.
- **Accountability.** Without an identifiable accountable party, the affected dimension stays at level 3 *(Controlled use)* or above, since safeguards need someone to apply them; the use can proceed only under external supervision.
- **Reversibility** works differently. A lasting effect is not harmful in itself, as when the benefit of training persists, so it sets no floor. But when harm cannot be undone, in the person or in data already used to train models, a calibration error costs more. Lowering the affected governance dimension then requires strong evidence, and its measures must prevent harm rather than rely on repairing it. Where data cannot be withdrawn from a trained model, as is often the case today, *prevention* carries the weight: for instance, training only on data people agreed to share for that purpose, or on de-identified data.

Table 2.
Weighing risk and opportunity in each governance dimension

	Risk	Opportunity
EMBODIED	What can realistically be inferred and how easily people can be re-identified[38], alone or combined with other bodily data[39]	Validated clinical, scientific, or personal value of what the signal measures
EMBEDDED	Power imbalances in the setting, vulnerability of those affected, and the cost of refusing	Benefits that reach those who bear the risk, compared with the best alternative
ENACTIVE	Changes in the person or the system over time, such as habits, self-understanding, or drift in adaptive systems	Benefits that persist with sustained use in real-world conditions
EXTENDED	Flows to third parties, linkage with digital data, model training, and outputs that act back on the person	Documented gains from sharing or linking data, such as earlier biomarkers or better-targeted treatment

The governance label follows the highest dimension, never an average or a sum, but the response covers all four. In **Figure 1**, the Embedded and Enactive dimensions make the use *Controlled*, yet Embodied and Extended still need the safeguards of their own level. When measures from different dimensions clash, the Scale proposes a simple rule: the measure of a dimension that sets the label prevails over that of a lower one, which applies as far as that allows. In the figure, Enactive calls for tracking how the use affects people over time, which means keeping data, while Embodied calls for collecting no more data than the use requires. Since Enactive is one of the two dimensions that set the label, tracking goes ahead, but only with the data it actually needs.

When dimensions that share the highest level clash, the assessment first looks for a measure that satisfies both and, failing that, favors the one that better protects those who bear the risk, recording why. If the use in **Figure 1** took place at work, for instance, Embedded might call for keeping individual data away from the employer, while Enactive would call for following each person's response over time. An independent party could then do the follow-up and give the employer only aggregate results.



For neurotechnology,
govern the use, not the source.

4.3 RECALIBRATION AND REVIEW

Every governance label and profile is reviewed when new evidence, a new context, or a change in scale of deployment triggers it, and these triggers should be tied to observable signs that a *specific use* is spreading or changing—such as investment in it, related patents, or new data infrastructures it relies on—so that review does not rest on discretion. This review regime appears with a solid outline in **Figure 1**. Provisional calibrations, including every Controlled use and any recalibration that depends on conditions that may lapse, also carry a scheduled review date and appear with a dashed outline in the figure.

A dimension moves when the use changes or new evidence arrives, not when its measures are met, since meeting them is what the level requires.

Five principles decide when new evidence moves a dimension; the second, third, and fifth adapt Policy Evidence Readiness Levels (PERL)[40], a framework for matching policy decisions to the maturity of the evidence behind them:

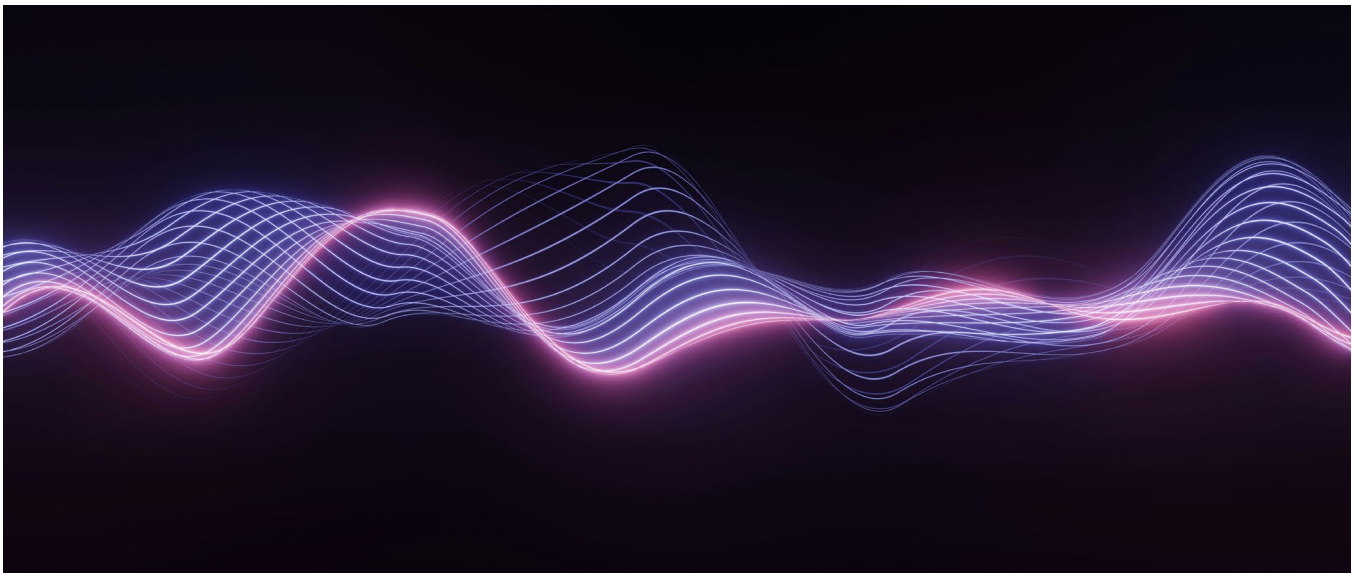
01 | Relevance.
A dimension moves only on new evidence about what placed it at its level: a proof that a device measures more accurately says nothing about whether its wearers can refuse.

02 | Maturity.
A single promising pilot justifies further testing; only results that hold across repeated studies and real-world conditions justify lowering the dimension so that the use can be extended to more people and to comparable settings.

03 | Context.
Evidence from one setting does not automatically apply to another: attention training shown to be safe for adult volunteers in a clinic is not necessarily safe for children in a classroom, so the classroom profile would move only with evidence from comparable settings. Evidence can also be shared: results from one deployment can support another in a comparable setting, which keeps the burden proportionate for smaller organizations.

04 | Asymmetry.
A credible sign of harm raises a dimension at once, because waiting costs more when people may be hurt; evidence of safety lowers it only when mature and from an equivalent setting.

05 | Burden of proof.
Those who deploy a use must show that its level can come down; regulators and users do not have to prove that it should stay where it is. The higher the level, the more independently that evidence should be checked.



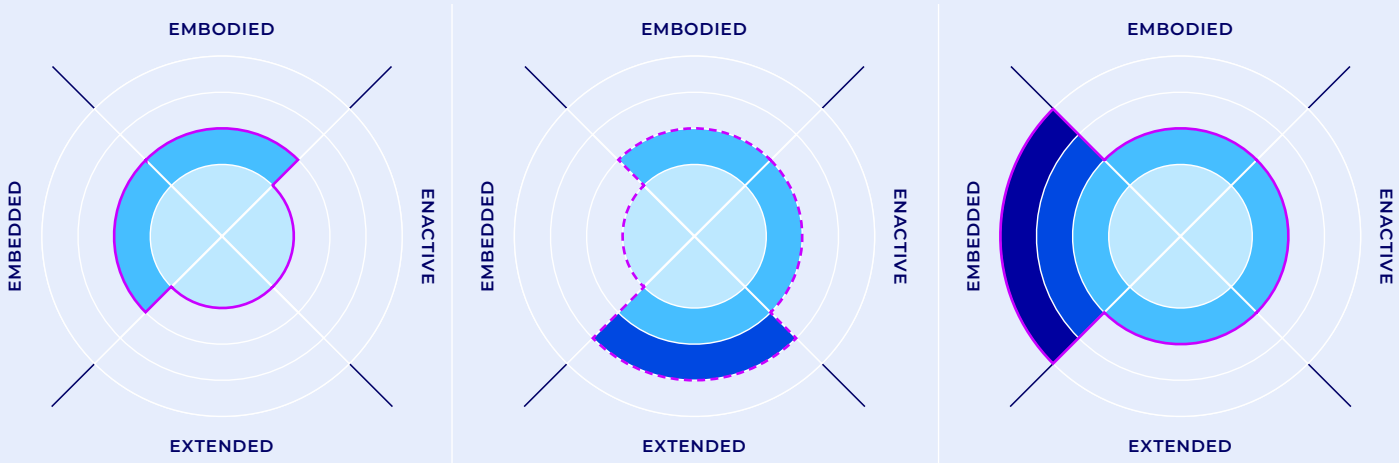


5. ILLUSTRATING THE SCALE: ONE HEADBAND, THREE USES

The three uses below do not map the field. They follow a single device across different settings to show how the Madrid Scale reads what changes from one use to another.

The uses are illustrative and the profiles indicative, as they do not assess any real product. A baseline helps: worn only for self-tracking, with data kept on the device, the headband raises no concern that existing rules do not already cover, and is Standard use (1-1-1-1). **Figure 2** shows how three other uses depart from it.

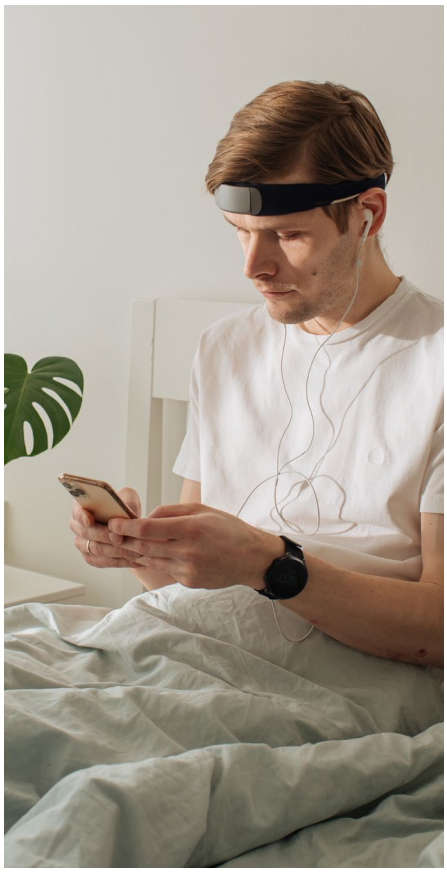
Figure 2.
The Madrid Scale Profile: Three illustrative uses



	(a) Flight deck	(b) Home	(c) Classroom
USE	Mandatory fatigue detection for pilots	App that merges neurodata with phone data to feed an AI assistant	Mandatory headband whose scores count toward grades
PROFILE	*Safeguarded use (2-2-1-1)*	*Controlled use (2-1-2-3)*	*Untenable use (2-4-2-2)*
REVIEW	Triggered (solid outline)	Triggered and scheduled, as a provisional calibration (dashed outline)	Triggered (solid outline)
SETS THE LABEL	Embodied (a neural inference, though narrow) and Embedded (required at work, though for safety)	Extended (what the app connects to)	Embedded: refusing would cost students, minors in a compulsory setting, part of their grade
WHAT WOULD MOVE IT	Nothing needed	Processing on the device, no model training: Safeguarded use (2-1-2-2)	Only a change of use. Voluntary training outside grading: *Controlled use (2-2-3-2)*, now led by Enactive (unproven effects on attention)

Source: Author's elaboration







Three points stand out. First, the device does not set the label; the use does, and the same headband ranges from *Standard* to *Untenable* depending on the use. Second, being required to wear it is not decisive in itself to set the label. On the flight deck **(Figure 2a)**, the inference is narrow, the purpose is safety, and the benefit reaches passengers as well as pilots and the airline. In the classroom **(Figure 2c)**, scores feed into grades, so refusing would cost students, who are minors in a setting they cannot leave, part of their grade. Third, what sets the label can lie far from the signal. At home **(Figure 2b)**, it is not what the headband records but what the app connects to: once brain data is merged with phone data to feed an AI assistant, the Extended dimension sets the label.

Each use can also move, in different ways. The home app **(Figure 2b)** moves through design: processing data on the device and keeping them out of model training changes what the app connects to, lowers the Extended dimension, and makes the use *Safeguarded*. The classroom use **(Figure 2c)** moves only by becoming a different use. Voluntary attention training, outside grading and with consent, becomes *Controlled*; the concern shifts to the Enactive dimension, as the training is meant to change how a child's attention works but evidence that such training helps children with attention difficulties is weak[41]. That result concerns this use, not schools or the headband per se.

The three uses also sit differently against existing rules. Without a declared medical purpose, none falls under the MDR[42]. The AI Act makes clear that detecting fatigue, as on a flight deck, is not emotion recognition, and its ban on inferring emotions in schools does not clearly reach attention scores; at most, scores that feed into grades could make the classroom system high-risk, which brings obligations but not a prohibition[43]. The profiles, however, do not depend on any of these laws. They describe the use, so an authority in another jurisdiction, from a US state to Chile, can read the same profile against its own rules.

This interoperability is one of the Madrid Scale's main strengths: it marks where existing rules suffice and where more may be needed, in terms that can travel across institutions and borders[44].

6. APPLYING THE MADRID SCALE: FIVE STEPS AND RECOMMENDATIONS FOR DECISION-MAKERS

The Madrid Scale applies to every use, clinical and research ones included. It is meant to be used by those who already make decisions, and to respond to the operational turn under way in neurotechnology governance[45]. It fits into what they already do, at design, funding, approval, procurement, and after deployment, rather than adding a new procedure.

Where an existing assessment already covers a dimension, such as a conformity assessment, a data protection impact assessment, or an ethics approval, it counts as evidence for that dimension rather than being repeated. Furthermore, because all read the same profile, an assessment made at one stage, by a funder or an ethics committee, can inform later ones, including regulatory review, instead of starting again each time.

6.1 APPLYING THE SCALE IN FIVE STEPS

01 | Define the use in its context: what the device does, where, by whom, on whom, for what purpose, and which systems it connects to.

02 | Set the level of each dimension, checking the twelve questions **(Table 1)** and weighing risk and opportunity **(Table 2)**.

03 | Read the profile: the highest dimension sets the label.

04 | Assemble the response, starting from the rules that already apply and adding the measures each dimension calls for.

05 | Record the assessment and set its review. The record should follow a fixed format: use and context, profile and label, measures, review regime (with its date, if scheduled), version of the Scale, and who carried out the assessment. That format will allow an assessment to be shared and compared.

The Madrid Scale is granular by design. An assessment can be as fine as the decision requires, down to one use in one setting for one group of people, and its measures go to the dimension that needs them. This **avoids one-size-fits-all rules**, whether blanket bans or blanket approvals: clearing one use of a device does not clear another.

Repurposing follows the same logic. When a device is used beyond its original purpose or setting, the new use does not inherit the earlier profile: it is a different use and needs its own assessment before it proceeds. That assessment, and the burden of proof, fall on whoever deploys the new use, such as an employer, a school, or an app provider, not on the people exposed to it. Existing law points the same way where it applies: under the AI Act, whoever changes a system's intended purpose so that it becomes high-risk takes on the provider's obligations[46], and the MDR does the same for anyone who changes the intended purpose of a medical device[47]. The Scale extends that logic to every use. Manufacturers of general-purpose devices, who may not know or disclose the final use, assess the uses they intend or promote and give deployers what they need to assess the rest, such as what the signal can reveal and where its data go. Developers can help from the design stage, by recording which uses they have assessed and which they have not, and regulators and buyers can ask for a profile whenever a device enters a new setting.

Assessors may disagree, and the Scale does not settle disagreement by authority. Each assessment records its reasons so that others can review and contest it, and the higher the level, the more independent that review should be (Section 4.3).

6.2 RECOMMENDATIONS FOR DECISION-MAKERS

The recommendations below are addressed to the main decision-makers who can apply the Madrid Scale.

Regulators: focus scrutiny where the profile points.
Regulators can use the profile to concentrate oversight on the dimension that sets a use's label, instead of applying the same requirements to every device of a kind. Recognizing recalibration matters as much: when a use changes in verified ways, a lower level should bring lighter obligations. Untenable use marks the point at which human rights limit calibration, and tools for review after deployment, still scarce outside clinical pathways, deserve investment[48].

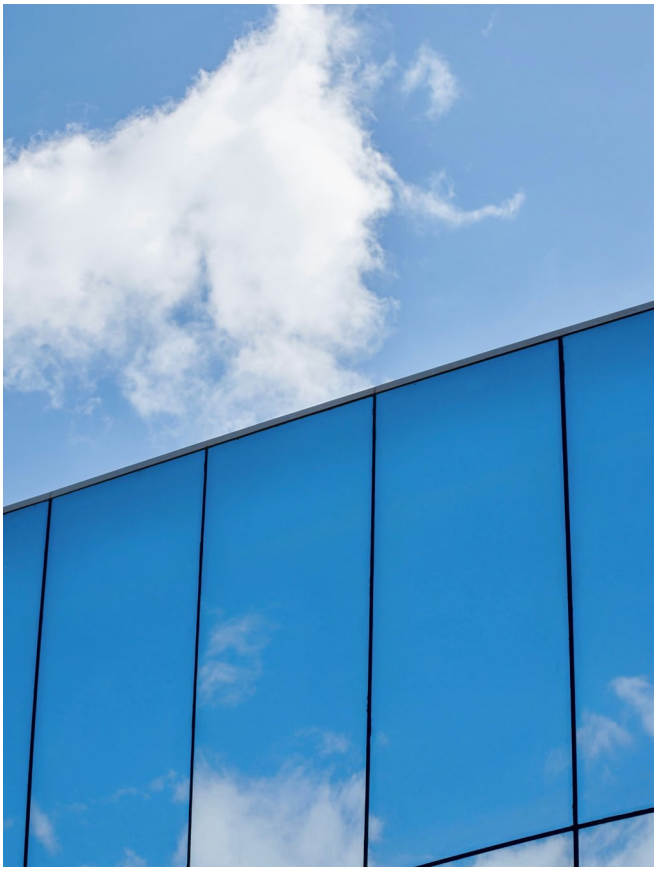
Public funders: scale requirements to the level and pay for evidence.
Funding calls can ask applicants for a profile and adjust their demands accordingly, so that low-level projects are not treated as high-risk ones, a burden that weighs most on smaller organizations[49]. For Controlled uses, funders are well placed to finance the evidence that will decide whether they can grow.

Ethics committees and public procurement: treat the profile as structured information.
The twelve questions (Table 1) give committees and buyers a shared checklist, and the profile a structured record that can be compared across proposals. Every assessment should state its review regime, and reviews of Controlled uses should include the people affected, who are often the first to notice harms that formal studies do not measure[50].

In neurotechnology, ethics and governance are a function, not a declaration: they need their own mandate, place, and resources.

Industry: assess from the design stage and claim no more than the evidence shows.
Companies can apply the Madrid Scale while a product is still being designed, when changing a use costs least. A documented low profile, open to review, signals trustworthiness to buyers, investors, and users, and claims should be backed by evidence proportionate to what they promise, as the earlier papers in The Neurotechnology Shift project series have urged[51].

International organizations: use the Scale as an interoperable reference.
OECD[52] and UNESCO[53] instruments set out principles and possible actions. The Madrid Scale offers a shared vocabulary that helps choose among them use by use, and because each profile describes the use rather than the law that governs it, jurisdictions with different rules can read the same profile in common terms. Its output follows a fixed format that can be recorded, compared across countries and over time, and tracked through indicators.



7. NEXT STEPS: BUILDING ON THIS VERSION

This version 1.0 sets out the logic of the Madrid Scale. Several elements need further work before it can be applied consistently and at scale, and they form the agenda for future versions.

Calibrated criteria.
The twelve questions **(Table 1)** and the weighing table **(Table 2)** guide judgment but do not yet define thresholds for each level in each dimension. A rubric with reference cases, tested for agreement between independent assessors, would make profiles more reproducible.

Roles along the value chain.
Manufacturers, deployers, and providers of data or AI services need clearly allocated roles, including for open-use platforms and for uses that cannot be disclosed.

Fit with existing assessments.
Future versions of the Scale should specify which parts of conformity assessments, data protection impact assessments, and ethics approvals count as evidence for each dimension, so that regulated uses are not assessed twice.

A proportionate path for small organizations.
Lowering a calibration level requires evidence that small companies may struggle to produce. A working definition of "comparable setting," shared evidence, and public funding could make that path proportionate.

Resources for review.
It remains to be defined who funds and performs reviews and independent verification.



8. FINAL REMARKS: GOVERN THE USE, NOT THE SOURCE

The Madrid Scale rests on a simple shift: govern the use, not the source. What matters is not only which device records a signal, what kind of data it produces, or what part of the brain it stimulates, but what it is used for, in which setting, on whom, and which ecosystem of data and AI systems it feeds.

Three clarifications follow:

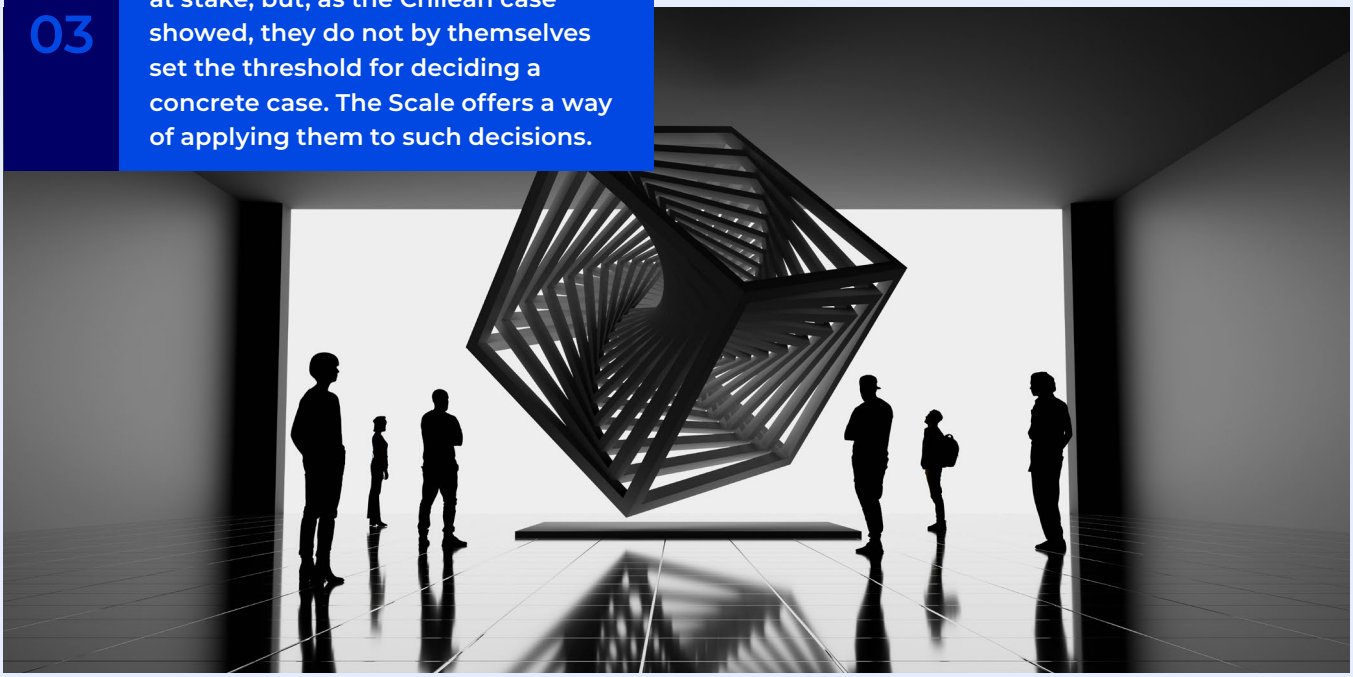
01 The Scale is not a classification of technologies, since it assesses one use in context at a time.

02 It is not a binding instrument or a list of prohibited products, and its outputs are illustrative and directional.

03 It does not replace rights. These say *why* protection matters and *what* is at stake, but, as the Chilean case showed, they do not by themselves set the threshold for deciding a concrete case. The Scale offers a way of applying them to such decisions.

Where a tool of this kind sits matters as much as what it contains. **Ethics and governance are a foundational function, not a declaration:** they settle *who* assesses a use, *where, how,* and *when*. **That function needs a mandate of its own**, a place in the institutional architecture, and resources to sustain it. It works best when built into design and data collection from the start, rather than bolted on as a final review. The Scale is meant to be used in that spirit, early in a project and again whenever a use changes.

Ultimately, neurotechnology governance does not have to choose between protecting people and enabling innovation. By calibrating each use to its context, the Madrid Scale aims to do both, so that beneficial uses can grow with confidence and those that cannot be justified are identified before they scale.



THE MADRID SCALE REFERENCE CARD

Four dimensions, four levels, twelve questions.

Protecting people while enabling innovation.

The Madrid Scale is an ecological, use-based framework for calibrating neurotechnology governance to risk and opportunity in context. It assesses uses, not devices: an application, in a given context, within the data ecosystem it connects to.

FOUR ECOLOGICAL DIMENSIONS OF NEUROTECHNOLOGY GOVERNANCE

Each weighs risk and opportunity within itself

EMBODIED

Neural signals are part of the body's biometric ecosystem: with heart rate, gaze, or voice, they support mental inferences.

Asks · What does the signal reveal, alone or with other bodily data?

EMBEDDED

The same device carries different risks and opportunities by setting, and weighs most on the vulnerable.

Asks · Where, by whom, on whom, and for what purpose?

ENACTIVE

Person and system shape each other through use, so effects keep emerging after deployment.

Asks · What changes, in the person and the system, as use continues?

EXTENDED

Neurodata feed digital and AI systems, and those systems act back on the person.

Asks · What does the use connect to, where do data and inferences go, and what returns?

FOUR CALIBRATION LEVELS

1 · STANDARD USE
Existing frameworks apply

2 · SAFEGUARDED USE
Added safeguards: standards, certification, limits on data use

3 · CONTROLLED USE
Pilot, sandbox, or monitoring while evidence is generated

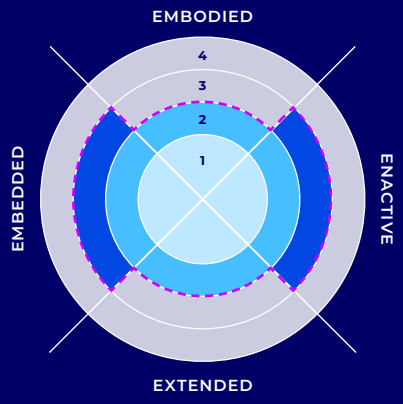
4 · UNTENABLE USE
Human rights set the limit: use not recommended; only a change of use can move it

READING THE MADRID SCALE PROFILE

BLOCKS
One per level reached, from light (1) to dark (4).

LABEL
The highest dimension sets it, never an average; the response still covers all four.

NOTATION
Label + profile (always Embodied-Embedded-Enactive-Extended). Here: *Controlled use (2-3-3-2).*



CLASHING MEASURES
The dimension that sets the label prevails. Between tied dimensions, seek a measure that satisfies both or, failing that, the one that better protects those at risk.

REVIEW REGIME
— Solid outline: reviewed when new evidence, a new context, or a new deployment scale triggers it.

– – Dashed outline: provisional calibration, also reviewed on a set date.

THREE CROSS-CUTTING CONDITIONS, IN EVERY DIMENSION

They can raise a dimension or make it harder to lower, never lower it

CONSENT
Does the person know what is happening, and can they genuinely refuse, as the use changes?

≥2 Partial consent

4 Consent fails altogether: covert inference, or a cost of refusal the person cannot bear.

ACCOUNTABILITY
Does someone identifiable answer for the use and offer a route for complaint?

≥3 No identifiable accountable party: the use proceeds only under external supervision.

REVERSIBILITY
Can harm be undone, in the person and in the data, including trained models?

No floor A lasting effect is not harmful in itself. But where harm could be permanent, a mistake costs more: lowering the dimension needs stronger evidence than usual, and safeguards must prevent harm, because it cannot be repaired.

© 2026 José M. Muñoz · V1.0 - October 2026 · CC BY-NC-SA 4.0

Based on: Muñoz J.M., The Madrid Scale for Neurotechnology Governance: A Use-Based Calibration of Risk and Opportunity in Context, IE CGC, October 2026
Cite as: Muñoz J.M., The Madrid Scale Reference Card, IE CGC, October 2026.

FIVE STEPS TO ASSESS A USE

1 **Define the use**
What the device does, where, by whom, on whom, for what purpose, and what it connects to

2 **Set each dimension's level**
Ask the twelve questions below; weigh risk and opportunity

3 **Read the profile**
The highest dimension sets the label (see overleaf)

4 **Assemble the response**
Rules that already apply (the level-1 baseline), plus the measures each dimension calls for at its level

5 **Record and set review**
In a fixed format (use, profile, measures, review regime, Madrid Scale version, assessor) that can be shared and compared

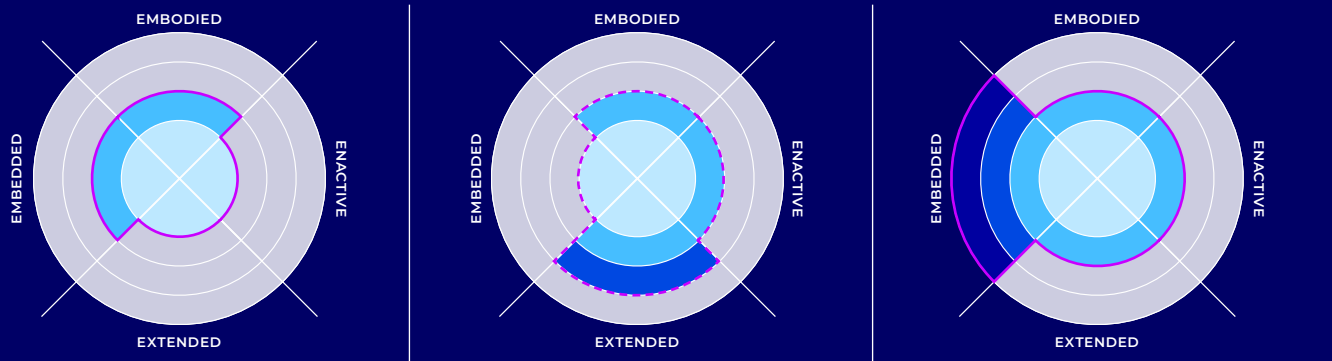
TWELVE QUESTIONS TO SET THE LEVEL OF EACH DIMENSION

The three conditions, asked in each dimension

	CONSENT	ACCOUNTABILITY	REVERSIBILITY
EMBODIED	Does the person know what is recorded and what can be inferred from it?	Who answers for the accuracy of what is inferred?	Can recorded data and inferences be deleted?
EMBEDDED	Can the person refuse without losing something they value, such as a job, a grade, or a service?	Who deploys the use, and where can those affected complain?	If the use feeds a decision about the person (e.g., a job, a grade, a service), can that decision be contested and reversed?
ENACTIVE	Is the person told how the system adapts and what it may change in them?	Who monitors effects over time and acts on them?	Do changes in the person fade when the use stops?
EXTENDED	Does the person know where the data go and agree to each further use?	Who answers for third parties that receive data or inferences?	Can shared data be recalled, including from trained models?

ILLUSTRATING THE MADRID SCALE · ONE EEG HEADBAND, THREE USES

Hypothetical uses; no real product is assessed here



FLIGHT DECK
Mandatory fatigue detection for pilots
Safeguarded use (2-2-1-1)
Review: triggered

HOME
App merging neurodata with phone data for an AI assistant
Controlled use (2-1-2-3)
Review: triggered and scheduled

CLASSROOM
Mandatory headband whose scores count toward grades
Untenable use (2-4-2-2)
Review: triggered

RECALIBRATION AND REVIEW

A dimension moves when the use changes or new evidence arrives. Five principles decide when new evidence moves a dimension:

- **Relevance** · Only new evidence about what placed a dimension at its level can move it.
- **Maturity** · A promising pilot justifies more testing; lowering a dimension needs results repeated in real-world conditions.
- **Context** · Evidence from one setting applies to another only if they are comparable.
- **Asymmetry** · A credible sign of harm raises a dimension at once; evidence of safety lowers it only when mature.
- **Burden of proof** · Those who deploy must show a level can come down; the higher the level, the more independent the check.

A new purpose or setting is a new use: a repurposed device does not inherit its earlier profile, and whoever deploys the new use assesses it and bears the burden of proof.

USING THE SCALE, BY DECISION-MAKER

- **Regulators** · Focus scrutiny on the dimension that sets the label; lighten obligations if justified by verified recalibration.
- **Public funders** · Ask for a profile, scale requirements to the level, and fund the evidence that Controlled uses need.
- **Ethics committees & procurement** · Use the twelve questions as a shared checklist and the profile as a comparable record.
- **Industry** · Assess from the design stage and claim no more than the evidence shows.
- **International organizations** · Use profiles as an interoperable reference, read against each jurisdiction's rules.

The Madrid Scale works best within a standing ethics and governance function, with a mandate, a place in the institutional architecture, and resources, built in from the design stage.

© 2026 José M. Muñoz · V1.0 - October 2026 · CC BY-NC-SA 4.0

REFERENCES

1 Irene Blázquez (Director, Center for the Governance of Change at IE University), Carlos Luca de Tena (Executive Director, Center for Governance of Change at IE University), Alberto Armengol (Senior Associate Director, Center for Governance of Change at IE University), Álvaro Matud (Institutional and Academic Director, Fundación Tatiana), Leticia Lombardero Vindel (Director of Planning and Development, Fundación Tatiana), José M. Muñoz (Research Director, The Neurotechnology Shift; Founding Director, INEGOV; Faculty Associate, Berkman Klein Center at Harvard University; Adjunct Professor of Neuroethics, University Carlos III of Madrid), Eva Ortega-Paíno (Secretary General for Research, Ministry of Science and Innovation of Spain), Francisco Javier Ponce Martínez (General Director, Sociedad Española para la Transformación Tecnológica), Carlos de la Cruz Molina (Director of Technology and Internationalization, Centro para el Desarrollo Tecnológico y la Innovación), Gloria Hernández (Head, Centro Nacional de Certificación de Productos Sanitarios), Liset M. de la Prida (Director, Cajal Neurosciences Center), Enrique J. Gómez (Director of the Neurotechnology MSc, Universidad Politécnica de Madrid), Elena Muñoz Marrón (Scientific Director, UNNE Instituto de Neuroestimulación; Director of the Neuropsychology MSc, Universitat Oberta de Catalunya), Santiago Canals (Deputy Director, Instituto de Neurociencias), Javier Díaz-Nido (Professor and Researcher, Universidad Autónoma de Madrid), Manuel Martín-Loeches (Professor, Universidad Complutense de Madrid), Aureli Soria-Frisch (Director of Neuroscience and R&D Neuroscience Manager, Starlab), Daniel Díez (CEO, Connectoma Neurotech; Associate Professor, IE University), Ignacio de Ramón Burgos (CEO, SincroLab), and José Piquer Martínez (Director, Fundación Neurocirugía, Educación y Desarrollo; Associate Professor, IE University).

2 Muñoz, J. M. (2026). Neurotechnological natives and the neurotechnology shift: A research agenda for the brain-as-interface era. *Annals of the New York Academy of Sciences, 1562*(1), Article e70381. https://doi.org/10.1111/nyas.70381

3 IE Center for the Governance of Change. (n.d.). *The neurotechnology shift. IE University.* https://www.ie.edu/cgc/research/neurotechnology-shift-cgc

4 European Group on Ethics in Science and New Technologies. (2026). *Governing neuro-AI: Towards an infrastructure approach*. Publications Office of the European Union. https://doi.org/10.2777/8619693

5 UNESCO. (2025). *Recommendation on the ethics of neurotechnology.* https://unesdoc.unesco.org/ark:/48223/pf0000394866

6 Management of Individuals' Neural Data Act of 2025, S. 2925, 119th Cong. (2025). https://www.congress.gov/bill/119th-congress/senate-bill/2925/text

7 European Brain Council (2026, October 1). *Trust as a competitive edge: The Seal of Responsible Neurotechnology gets ready for launch.* European Brain Council. https://www.braincouncil.eu/seal-of-responsible-neurotechnology-gets-ready-for-launch

8 International Center for Neuroscience and Ethics. (2026, February 16–17). *Shaping common ground in neurotech regulation and governance* [1st CINET Neurotechnology Conference]. Fundación Tatiana, Madrid, Spain. https://www.cinetcenter.com/en/event/neurotech-regulation-and-governance

9 BrainMind. (2026). *Asilomar for the Brain and Mind* [Summit]. Asilomar Conference Grounds, Pacific Grove, CA, United States. https://www.asilomarbrainmind.org

10 Hensing, J., & Schlecht, P. (2025). *Action potentials: Neurotechnology, brain-computer interfaces, and implications for Germany's and Europe's foreign & security policy.* Global Public Policy Institute. https://gppi.net/2025/08/11/action-potentials

11 Muñoz, J. M., et al. (2026). Neurotechnology governance: The operational turn. *Trends in Biotechnology.* https://doi.org/10.1016/j.tibtech.2026.09.016

12 OECD. (2019). *Recommendation of the Council on Responsible Innovation in Neurotechnology* (OECD/LEGAL/0457). https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0457

13 UNESCO, *Recommendation on the ethics of neurotechnology,* op. cit.

14 OECD. (2025). *Neurotechnology toolkit: To support policymakers in implementing the OECD Recommendation on Responsible Innovation in Neurotechnology.* https://www.oecd.org/content/dam/oecd/en/topics/policy-sub-issues/emerging-technologies/neurotech-toolkit.pdf

15 Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices. (2017). EUR-Lex. https://eur-lex.europa.eu/eli/reg/2017/745/oj

16 Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). (2016). EUR-Lex. https://eur-lex.europa.eu/eli/reg/2016/679/oj

17 Bernáez Timón, L. (2026). *The expanding reach of neurotechnology beyond the clinic: Into consumer markets and the workplace* (The Neurotechnology Shift, Paper #1). IE Center for the Governance of Change. https://static.ie.edu/CGC/IE_CGC_NeurotechologyShift_Neurotechnology_Beyond_The_Clinic.pdf

18 BrainMind. (n.d.). *BCI universe*. https://brainmind.org/bci

19 Salles, A., et al. (2024). *Towards inclusive EU governance of neurotechnologies.* Institute of Neuroethics & Centre for Future Generations. https://cfg.eu/towards-inclusive-eu-governance-of-neurotechnologies

20 Schultz, S. R. (2026). *Translating medical neurotechnologies into non-clinical contexts* (The Neurotechnology Shift, Paper #2). IE Center for the Governance of Change. https://static.ie.edu/CGC/IE_CGC_NeurotechologyShift_Report%232_TranslatingMedicalNeurotechnologies.pdf

21 Ruiz, S., et al. (2024). Neurorights in the Constitution: From neurotechnology to ethics and politics. Philosophical Transactions of the Royal Society B: Biological Sciences, 379(1915), Article 20230098. https://doi.org/10.1098/rstb.2023.0098

22 Muñoz, J. M., et al. (2024). Effects of the first successful lawsuit against a consumer neurotechnology company for violating brain data privacy. Nature Biotechnology, 42(7), 1015–1016. https://doi.org/10.1038/s41587-024-02303-2

23 Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). (2024). EUR-Lex. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

24 Crewe, M. (2026, August 24). *California moves to regulate workplace neurotech. Neurofounders.* https://www.neurofounders.co/articles/california-moves-to-regulate-workplace-neurotech

25 Gallagher, S. (2023). *Embodied and enactive approaches to cognition. Cambridge University Press.* https://doi.org/10.1017/9781009209793

26 Magee, P., et al. (2024). Beyond neural data: Cognitive biometrics and mental privacy. *Neuron, 112*(18), 3017–3028. https://doi.org/10.1016/j.neuron.2024.09.004

27 Hendriks, S., et al. (2026). Risk-calibrated sharing of human brain data. *Brain.* Advance online publication. https://doi.org/10.1093/brain/awag205

28 Judel, Á., et al. (2026). Turning biometrics into a commodity raises privacy and governance risks. *Nature Human Behaviour.* Advance online publication. https://doi.org/10.1038/s41562-026-02587-1

29 Wang, S., et al. (2023). Closed-loop adaptive deep brain stimulation in Parkinson's disease: Procedures to achieve it and future perspectives. *Journal of Parkinson's Disease, 13*(4), 453–471. https://doi.org/10.3233/JPD-225053

30 Muñoz, Neurotechnological natives and the neurotechnology shift: A research agenda for the brain-as-interface era, op. cit.

31 European Group on Ethics in Science and New Technologies, *Governing neuro-AI: Towards an infrastructure approach,* op. cit.

32 Hendriks et al., *Risk-calibrated sharing of human brain data*, op. cit.

33 Gallagher, *Embodied and enactive approaches to cognition*, op. cit.

34 Kaye, J., et al. (2015). Dynamic consent: A patient interface for twenty-first century research networks. *European Journal of Human Genetics, 23*(2), 141–146. https://doi.org/10.1038/ejhg.2014.71

35 Muñoz et al., Neurotechnology governance: The operational turn, op. cit.

36 Hendriks et al., Risk-calibrated sharing of human brain data, op. cit.

37 Muñoz, et al. Neurotechnology governance: The operational turn, op. cit.

38 Hendriks et al., Risk-calibrated sharing of human brain data, op. cit.

39 Magee et al., Beyond neural data: Cognitive biometrics and mental privacy, op. cit.

40 Cotton, C. (2026). Rigor is not readiness: The PERL framework for evidence-based policy. *Science, 393*(6810), 463–465. https://doi.org/10.1126/science.aef3529

41 Westwood, S. J., et al. (2025). Neurofeedback for attention-deficit/hyperactivity disorder: A systematic review and meta-analysis. *JAMA Psychiatry,* 82(2), 118–129. https://doi.org/10.1001/jamapsychiatry.2024.3702

42 Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, op. cit.

43 Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), op. cit.

44 Muñoz et al., Neurotechnology governance: The operational turn, op. cit.

45 Ibid.

46 Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), op. cit.

47 Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, op. cit.

48 Muñoz et al., Neurotechnology governance: The operational turn, op. cit.

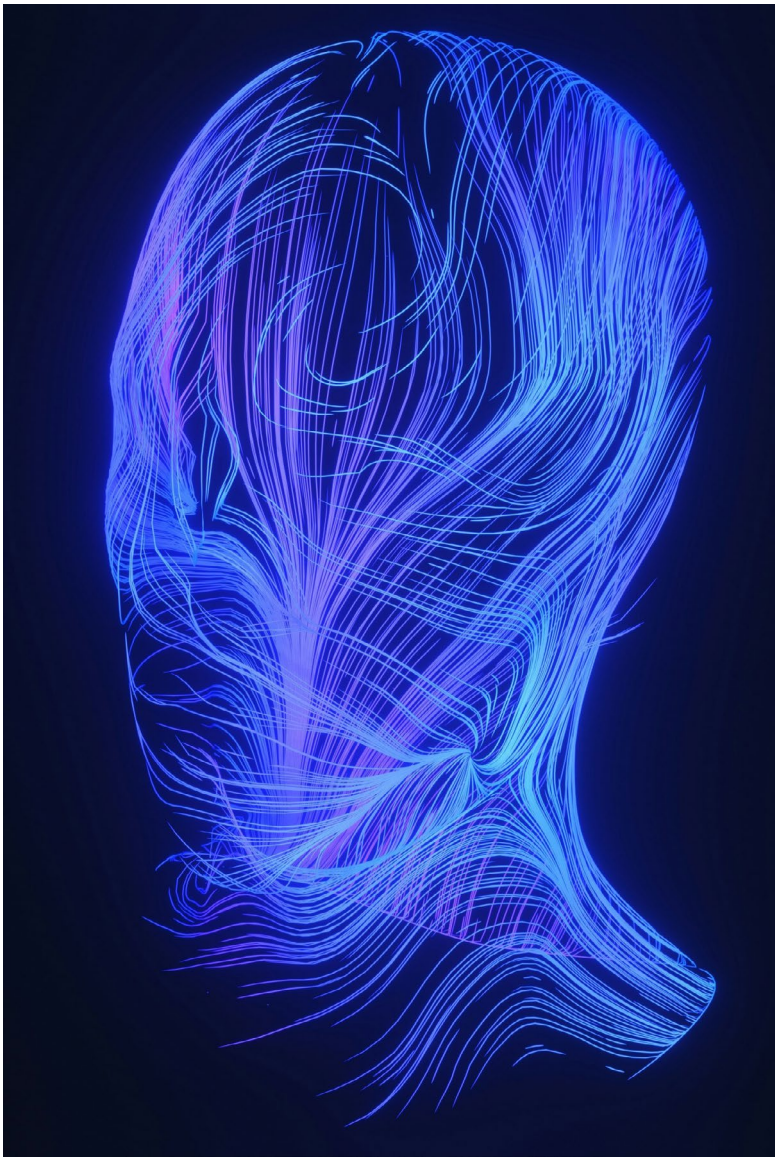
49 Ibid.

50 Ibid.

51 IE Center for the Governance of Change, The neurotechnology shift, op. cit.

52 OECD, *Recommendation of the Council on Responsible Innovation in Neurotechnology*, op. cit.

53 UNESCO, *Recommendation on the ethics of neurotechnology,* op. cit.



IE UNIVERSITY
CENTER FOR THE GOVERNANCE OF CHANGE

AUTHOR:
José M. Muñoz

RECOMMENDED CITATION:
Muñoz J. M., *The Madrid Scale for Neurotechnology Governance: A Use-Based Calibration of Risk and Opportunity in Context*, IE CGC, October 2026

© 2026 José M. Muñoz. Published by CGC, Madrid, Spain
Design: studioazola.com
Images: Shutterstock, Unsplash

This work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License. To view a copy of the license, visit creativecommons.org/licenses/by-nc-sa/4.0

cgc.ie.edu