NOVA Policy Network
Policy Brief

# Operationalizing the OECD Guidelines for Responsible Neurotechnology

*Evaluating how member countries are implementing the 2023 OECD Recommendation on Responsible Neurotechnology*

<Murari Ambati>

Policy Brief

Date, October 21, 2024

# Executive Summary

In 2023, the Organisation for Economic Co-operation and Development (OECD) approved the Recommendation on Responsible Neurotechnology, the first policy guidance in the world aimed at guiding the ethical design and deployment of brain-related technologies. The Recommendation—endorsed by 38 member states—sets out eight essential principles that include autonomy, beneficence, transparency, proportionality, mental integrity, inclusivity, accountability, and responsible innovation. It is intended to be relevant across the entire range of neurotechnologies, from medical-grade brain-computer interfaces (BCIs) to non-clinical cognitive enhancement technologies and commercial neurodata products. The Recommendation is forward-looking, anticipating complex ethical risks related to mental privacy, dual-use potential, and the integration of AI systems and neural interfaces.

However, the Recommendation is legally non-binding and contains no reporting or enforcement, relying exclusively on voluntary adoption by national governments and other stakeholders as appropriate. This decision, while diplomatically expedient, has resulted in very patchy implementation across jurisdictions. Some countries—such as Germany and Canada—have gestured towards alignment with either their existing biomedical ethics framework or AI principles. Some, like Japan and the United States, have yet to make concrete regulatory mechanisms implementing the OECD's basic neurotechnology ethics. Implementation in nearly all cases is piecemeal, with no states taking comprehensive measures for integrating the OECD's values into their domestic legal, technological, or research ecosystems.

This gap between global recommendation and domestic application entails enormous risks. Without legal definitions, mental integrity remains an abstract notion. Without sectoral norms, consumer neurotechnology markets expand in regulatory vacuums. Without specialized neuroethical review processes, new brain-interface technologies may bypass essential human rights assessment. And without harmonized international oversight, the Recommendation itself may not garner more than symbolic endorsement.

This note evaluates the current level of implementation of the Recommendation across OECD countries, identifies the structural and conceptual bottlenecks to operationalization, and proposes specific policy instruments—like National Implementation Charters, performance-based compliance metrics, and a multilateral Neurotechnology Oversight Forum. Unless and until countries move from normative alignment to policy action, the Recommendation risks becoming an ethical high-water mark with no tangible effect. In order for neurotechnology to go forward in a way that is compatible with democratic values and cognitive rights, governance must evolve from aspirational principles to actionable architecture.

# Introduction

The exponential growth of neurotechnology—from non-invasive neural monitoring systems and invasive brain–computer interfaces to neuropharmaceuticals and AI-driven neurodata analytics—has ushered in a new era of collaboration between technology and the human mind. The potential of these technologies in medicine, communication, education, and human performance is vast. Yet they also raise urgent questions about privacy, agency, cognitive liberty, and the adequacy of current legal frameworks to govern tools with direct effects on mental states. In response to these increasing demands, the Organisation for Economic Co-operation and Development (OECD) released its 2023 Recommendation on Responsible Neurotechnology, the first internationally agreed upon framework solely dedicated to the ethical governance of brain-related technologies.

The Recommendation is a consolidation of interdisciplinary stakeholder consultations and formulates eight normative principles for the responsible governance, development, and deployment of neurotechnologies. These span ensuring mental integrity to safeguarding user autonomy and informed consent, promoting inclusivity and benefit-sharing, and establishing transparency and accountability mechanisms. Aspirational and internationally applicable as the Recommendation is, it is by definition voluntary. Unlike binding treaties or regulatory standards, it relies on the goodwill of OECD member governments to convert its recommendations into national policy agreements, legal documents, institutional review procedures, and industrial norms. It is not accompanied by enforcement mechanisms, reporting obligations, or compliance metrics.

This voluntary system is both a promise and a danger. On the one hand, the Recommendation provides a flexible ethical paradigm that can be translated by states to fit their legal traditions and sociopolitical circumstances. On the other hand, it creates a governance gap in which states can symbolically ratify the Recommendation without substantively following its principles. Lacking operational clarity, definitional specificity, or formal oversight, the Recommendation risks being aspirational rather than actionable—especially in fast-moving areas such as consumer neurotechnology and AI-neurointegration.

This policy brief addresses that challenge. It surveys how some OECD countries are interpreting—or disregarding—the Recommendation, identifies structural gaps in translation from principle to policy, and proposes a pragmatic roadmap to transform the Recommendation into a living governance architecture. By exploring the implementation lag, this brief underscores the need to move from normative consensus to institutional design to safeguard cognitive rights and assure neurotechnologies evolve within institutions that respect human dignity, justice, and democratic accountability.

# Issues / Policy Gaps

Although the OECD made a major step towards creating a global ethical code for neurotechnology, its application as a national practice has evidenced various conceptual and structural flaws that limit its functional efficacy. Most crucially among them is the absence of legally binding force. The fact that the Recommendation is recommended, rather than obligatory, renders it susceptible to politics' and institutional interpretation's will, which has led to huge discrepancies as to either how or even whether nations engage with its principles. In practice, it has provided a governance gap on which neurotechnology can rapidly develop and implant itself with minor or intermittent ethical regulation across boundaries.

One of the deepest lacunae is in the definitional uncertainty that hangs around decisive concepts such as "mental integrity" and "cognitive liberty." While both are rich in ethical substance at the conceptual level, they lack legal moorings in the majority of legal systems. Policymakers and regulatory authorities thus do not have working tools to decide whether a neurotechnology violates or respects these concepts. Without definitional certainty, enforcement becomes impossible, and ethical principles become rhetorical tropes rather than solid guidelines.

A second major issue is a mismatch between the Recommendation's ethical needs and existing national regulatory systems. Most countries apply neurotechnologies to frameworks of general biomedical research, medical device law, or data protection regimes—none of which were designed with neural interfaces or brain data in mind. For example, neural data falls outside the scope of current data protection legislation such as the GDPR that does not consider it a special or sensitive category unless it can be explicitly associated with health diagnostics. This leaves specific regulatory blind spots, particularly given that commercial actors begin to collect and commodify brain-focused data outside of the clinic.

The third policy gap arises due to the institutional limitations of ethics review processes. Institutional Review Boards (IRBs) or Research Ethics Committees (RECs) in most nations do not have a representative with expertise in neuroscience, neuroethics, or cognitive science. The institutions cannot therefore evaluate the whole array of dangers pertaining to advanced neurotechnologies, particularly those employing AI models for inferring or altering mental states. Ethical assessment, as a first-line defense against harmful or coercive processes, is thus rendered incoherent and perhaps ineffective.

Another critical issue is the absence of any monitoring procedure among OECD states to track adoption, implementation, or impacts. Unlike designs involving peer-review mechanisms or benchmarking instruments, the Recommendation does not mandate reporting of progress or aligning policies with standardized aims. The resulting absence of transparency makes it difficult to judge what states are indeed implementing the principles and what states are not, ultimately weakening collective responsibility.

Finally, civil society involvement is anemic across the board. Although the Recommendation values much inclusivity, efforts at participation have been sporadic in most nations to engage neurodivergent individuals, patient groups, or underrepresented groups in the shaping of national neurotech policy. This weakens the democratic legitimacy of any regime of regulation that purports to safeguard mental autonomy and equitable innovation. The result is an increasingly widening gap between policy practice and ethical intention—a gap that subjects important aspects of the human mind to increasingly greater technological intrusion without concomitant protection.

# Policy Recommendations

To overcome the operational shortfalls of the 2023 OECD Recommendation on Responsible Neurotechnology, this brief proposes a multi-level governance framework that translates normative values into enforceable, quantifiable, and jurisdiction-limited regulatory instruments. The eventual objective is to create scalable policy infrastructure that can both respect national diversity and ensure minimal global benchmarks for cognitive rights, ethical innovation, and mental sovereignty.

The first of these is the development of national \textit{Neurotechnology Implementation Charters} (NICs), released formally by OECD countries. They would be a statutory or executive policy tool that integrates the eight OECD principles into domestic law, research funding rules, public procurement, and regulatory review channels. These charters should operationally clarify terms such as "mental integrity" and "neurodata," and clearly indicate the agencies and oversight bodies that will enforce them. Significantly, the NICs would distinguish clinical, commercial, and military applications of neurotechnology, scaling models of governance to respective sector-specific risks.

The second is the establishment of a centralized \textit{OECD Compliance Architecture Framework} (OCAF), which would supply standardized instruments, audit guidelines, and ethics checklists to be utilized by member states. It would promote the adoption of neuroethical review procedures in research ethics committees, templates for informed consent in neural data, and dual-use risk assessment metrics in national funding calls. It would also recommend that each member state have a Neurotechnology Regulatory Coordinator, whose task would be to promote cross-ministerial alignment between science, health, data protection, and innovation authorities.

To quantify and benchmark implementation across OECD members, this brief introduces a mathematical model of policy alignment, referred to as the *Neurotechnology Implementation Index* (NII). Let each country $i$ be evaluated across the eight OECD principles $j = 1, 2, \ldots, 8$, where each compliance level $c_{ij} \in [0, 1]$ is normalized to a unit scale. Let $w_j \in [0, 1]$ be the weight assigned to each principle, reflecting its global ethical priority. The overall score for country $i$ is computed as:

$$NII_i = \sum_{j=1}^{8} w_j \cdot c_{ij}$$

This index produces a value $NII_i \in [0, 1]$, with higher scores indicating deeper and more comprehensive alignment. The weights $w_j$ may be determined through expert consensus or ethical priority ranking conducted by the OECD Neurotech Forum. Countries with scores below a minimum implementation threshold—say, $NII_i < 0.60$—would be subject to targeted peer engagement or compliance assistance programs.

To visualize this model, the OECD should launch a *Global Neurotechnology Dashboard*, an open benchmarking platform with NII scores plotted across member states. This would not only facilitate transparency and accountability but induce proactive adoption of best practice. In addition, involving civil society organizations and neuroethics experts in the OECD's governance framework through a standing *Multilateral Neurotechnology Oversight Forum* would ensure participatory legitimacy and ongoing scrutiny of policy.

Lastly, the OECD Recommendation's transformation into a living system of governance hinges on the intersection of national legal intervention, international coordination, quantifiable compli-

ance modeling, and democratic participation. It is only by embedding its norms into actual institutional designs that the OECD can work loyally to its mandate to guide neurotechnology development along a trajectory according to justice, security, and cognitive dignity.

# Conclusion

The OECD Recommendation on Responsible Neurotechnology is a bold and long-overdue effort to guide the ethical trajectory of an increasingly dynamic technology. It establishes a normative vision for protecting mental integrity, advancing cognitive liberty, and ensuring brain-linked innovation aligns with democratic principles. But as this brief has shown, the gap between promise and practice remains wide. Without concrete legal instruments, oversight processes, or institutional capacity, the Recommendation runs the risk of being a sincerely intended but essentially empty policy gesture—particularly in light of accelerating the release of neurotechnologies that influence thinking, identity, and action by commercial and state entities.

Closing the gap requires work. Governments at the national level must move beyond token expressions of support and establish binding neurotechnology implementation charters that translate ethical principles into practical policy. International institutions, even the OECD, must facilitate this through the provision of harmonized compliance architectures, audit templates, and ethics assessment tools tailored to the particularities of neural interfaces, neurodata, and dual-use systems. Not less crucial is the adoption of a single, core performance metric—e.g., the Neurotechnology Implementation Index—that enables benchmarking, transparency, and cross-country learning. It is only by keeping such tools in reserve that we can ensure that the Recommendation evolves from static report to dynamic governance structure.

On a deeper level, healthy neurotechnology governance is more than just regulation—protecting the margins of the human brain in the age of deeper machine embedment. The technologies in question have the power not only to cure and enhance, but to manipulate, commodify, and even exploit our most intimate neural processes. The ethical stakes are therefore not academic but existentially gravely serious. The OECD Recommendation setting out authoritatively for the first time a global framework, has enormous potential—but it will only be activated if states are willing to act passionately, openly, and inclusively. The future of cognitive rights depends on it.