N✦VA

NOVA Policy Network
Policy Brief

Evaluating the Ethical Governance Gaps in the NIH Brain Initiative

Addressing Oversight, Consent, and Dual-Use Risks in America's Flagship Neurotechnology Program

<Murari Ambati>

Policy Brief

Date, August 15, 2024

Executive Summary

The US BRAIN (Brain Research through Advancing Innovative Neurotechnologies) Initiative is the most ambitious federally funded neuroscience project ever. Supported initially in 2013 and significantly amplified through BRAIN 2.0 in 2020, it has invested over $3 billion in state-of-the-art neurotechnological tools for brain mapping, neural modulation, and cognitive enhancement. But the Initiative's frenetic pace of scientific advance has left it behind in its ethical shield.

This policy brief discovers critical governance gaps in the BRAIN Initiative's current regulatory scheme. They are: (1) inadequate enforcement provisions for ethical review, (2) no uniform consent templates for neural experimentation, (3) no dual-use risk disclosure requirements in grant proposals, and (4) inadequate federal guidance on ownership, protection, and utilization of neural data. Without strong action, these gaps have the potential to lead to undesirable consequences ranging from cognitive coercion to unethical surveillance and biased access to neurotechnologies.

To achieve this, this memo recommends the creation of a Federally Mandated Neuroethics Review Board (FNRB), mandatory dual-use risk assessments of NIH-funded research, the development of tiered informed consent procedures tailored to the levels of neurotechnology risk, and the institutionalization of a national neurodata governance system.

As neurotechnological science becomes more tightly entwined with national security, health equity, and cognitive rights, strong ethical governance is in demand. The BRAIN Initiative needs to become not just a scientific leader, but also a model of policy vision and responsible innovation.

Introduction

The human brain can be best described as the most complex system in the known universe. In recognition of the brain's central role in determining individual and public health, the United States launched the BRAIN (Brain Research through Advancing Innovative Neurotechnologies) Initiative in 2013. The Initiative, spearheaded by the National Institutes of Health (NIH) and a cooperating set of federal and private entities, aimed to accelerate the development of new technologies to chart the brain's structure and activity. With the 2020 BRAIN 2.0 expansion, the program expanded exponentially—integrating investments in large-scale brain-computer interfaces (BCIs), live capture of neural data, circuit-level manipulations, and experimental cognitive enhancement methods.

To date through 2025, over $3 billion has been invested in the Initiative, making it one of the most ambitious neuroscience initiatives in the world. While the scientific breakthroughs are obviously encouraging—everything from novel diagnostic for mental illness to early neural prosthetics—the BRAIN Initiative has also triggered fundamental questions about ethical control, subject autonomy, neurodata privacy, and dual-use harms. These are particularly urgent considerations considering the Initiative's impact on global standards for the development of cognitive technology.

Despite the publication of NIH's *BRAIN Neuroethics Roadmap* in 2020, the Initiative's regulatory controls on ethics remain advisory and not enforceable, contrary to the high-stakes nature of the research it supports. The absence of robust regulatory scaffolding leaves enormous vulnerabilities: players may not be sufficiently well-informed of the long-term effects of neurotechnological treatments; information based on brain activity may not be sufficiently well-covered by existing laws such as HIPAA; and technologies initially framed for healthcare may be exceptionally readily redirected towards military, corporate, or surveillance goals.

Secondly, the process through which NIH financing is distributed in BRAIN also lacks open criteria for assessing the ethical and social impacts of research proposals. While scientific discovery is given priority, sufficient attention is not offered towards risk avoidance, dual-use utilization, and equity of access and distribution. In the absence of ethical equivalence, the Initiative has the potential to entrench existing societal inequalities and support coercive regimes of the mind, particularly in militarized or work environments.

This policy brief critically examines the BRAIN Initiative's governance framework from four main points: (1) the strengths and limitations of current ethical oversight, (2) the risk of dual-use neurotechnology, (3) the adequacy of informed consent processes, and (4) the regulatory loopholes in neural data privacy. By examining these deficits, the brief proposes a package of policy reforms that can strengthen the Initiative's ethical moorings without constraining its scientific potential.

Lastly, the purpose of this analysis is not to halt innovation, but to ensure that the great promise of brain science is being balanced by an equal fervent commitment to cognitive liberty, openness, and democratic responsibility.

Governance Challenges in the BRAIN Initiative

The promise of neurotechnological advancement, while deep, is offset by a difficult set of unresolved moral, legal, and political issues that are contained within the governance structure of the BRAIN Initiative. These are not afterthoughts—they are intrinsic to the design of the Initiative. The following sections identify four conceptual gaps that policy must address in the near future: (1) decentralized ethical oversight frameworks, (2) ubiquitous dual-use infrastructures, (3) conceptual incoherence in informed consent practices, and (4) illegality of neurodata.

Fragmented and Decentralized Ethical Oversight

The BRAIN Initiative is a decentralized funding system under which ethical review is outsourced to local Institutional Review Boards (IRBs) within individual grantee institutions. IRBs are mandated by the Common Rule, but their neuroethical expertise varies greatly. There is no single federal regime under which there is an assurance that IRBs possess adequate domain-specific expertise in the field of cognitive modulation, BCIs, or neural data analytics.

This decentralization generates a conceptual void: ethical assessment becomes procedural rather than substantive. Protocols are approved for "compliance" rather than epistemic integrity or pre-emptive harm analysis. Additionally, the NIH's *Neuroethics Roadmap*—visionary as it is—lacks enforceability and is not binding. The absence of a central neuroethics adjudicating body generates a governance vacuum within which ethically dubious projects can pass by overcoming mere compliance barriers.

Latent Dual-Use Architectures and Strategic Silence

Much of the innovation enabled by the BRAIN Initiative is dual-use by nature—capable of both civilian and military or coercive application. Such examples include targeted neuromodulation, real-time mood monitoring, or remote control of prosthetic limbs. However, NIH funding programs have no dual-use disclosure requirement. Researchers do not need to assess or report on the ways in which their technologies can be repurposed into non-consensual behavioral control, psychological warfare, or surveillance.

This structural exclusion is a form of "strategic silence," wherein dual-use capability is left out not due to ignorance, but simply because there is no formalized epistemic space wherein it must be evaluated. Whereas biological or nuclear research fall under the jurisdiction of DURC (Dual-Use Research of Concern) policy, neuroscience is de facto exempt from dual-use review. This asymmetry is ethically intolerable, particularly inasmuch as cognitive liberty is at the very foundation of democratic self-governance.

Conceptual Incoherence in Informed Consent Frameworks

Informed consent continues to form the basis of research ethics, but in neurotechnology it is under tension in its ideals. Participants in BCI experiments, for example, are requested to provide

consent to interventions whose mechanism of action is not fully understood and whose long-term neuropsychological effects are unknowable. Consent in such conditions is pseudovoluntary: it exists for bureaucratic form but lacks epistemic content.

Second, traditional models of consent postulate essential identities and wise foresight. The opposite is that neurotechnological interventions could alter cognition, affect, memory, or will - leading to the possibility of retroactive or altered preference patterns. When a neurodevice alters the motivational landscape of a subject, can prior consent persist? No direction is available from existing governance frameworks.

This is compounded in high-risk populations, such as adolescents or the mentally ill, where cognitive and legal capacity for consent is already at risk. The lack of tiered consent processes, graduated by the severity of intervention, harms both autonomy and ethical integrity.

Legal Invisibility of Neurodata and Cognitive Signatures

Most in urgent need of regulation, perhaps, is the processing of neural data. Brain signals - whethererer collected by EEG, fMRI, or invasimplants - constitutitute an emergingemergingemerging class of information: cognitive signatures. These data are biologically distinct, affectively dense, and inferentially strong. But they are in a legal gray area, excluded from HIPAA if collected outside the clinic, and not clearly within the ambit of biometric data protections in GDPR or U.S. state law.

Neurodata is not health data or even behavioral metadata—it is a hybrid entity, useful for reconstructing internal states, identifying individuals, or inferring predispositions. The failure to define neurodata as a legal category has resulted in its passive commodification. Without a statutory regime of neurodata, corporations and researchers could be legally permitted to store, trade, and monetize neural signatures without participant awareness or control.

This is not a theoretical issue: the increasing commercialization of transportable EEG equipment, consumer-grade BCIs, and neuroenhancement programs indicates cognitive monitoring is arriving in the public marketplace. The BRAIN Initiative, by not grappling with these legal ambiguities, could be putting itself at risk of condoning a world in which neural privacy is not a right, but a transaction.

Policy Recommendations

Effective management of the BRAIN Initiative goes beyond reactive oversight—it requires an active ethic of architecture built into the design and implementation of federal neuroscience support. Closing the Initiative's foundation gaps involves synthesizing ethical theory, legal systems, computational modeling, and cross-regulatory science. Here is a multi-dimensional governance framework founded on cognitive autonomy, anticipatory regulation, and epistemic humility.

A federally required national centralized review infrastructure will have to be established to replace the current decentralized and variable Institutional Review Board (IRB) system. This agency—temporarily called the Federally Mandated Neuroethics Review Board (FNRB)—would be a specialist review and accreditation body within the NIH with interdisciplinary staff in neuroethics, computational neuroscience, regulatory law, and biosecurity. Each project funded by the BRAIN Initiative would be independently audited by FNRB prior to expenditure of funds. Of key significance, the FNRB would have investigatory powers and make binding decisions as to the ethics acceptability of cognitive modulation research, neural implants, and behavioral inference technologies.

In parallel, a standard dual-use risk categorization system should be integrated into NIH funding policies. A research proposal should include a Dual-Use Disclosure Form (DUDF), designating systematically potential abuse situations and reasonable abuse routes. To simulate and order such dangers, we propose a continuous dual-use exposure function:

$$\mathcal{R}(t) = \alpha \cdot \left(\frac{\partial T}{\partial U} \right) \cdot e^{\beta D(t)}$$

where $\mathcal{R}(t)$ is the real-time risk exposure at time t, T is technological capability, U is intended use, $D(t)$ is diffusion into non-regulated domains (e.g., military or commercial sectors), α is a policy sensitivity coefficient, and β captures exponential increase in vulnerability as diffusion grows. This model may be used by NIH policy analysts to simulate exposure under various regulatory lag scenarios, allowing ethical triage and timeline prioritization.

A third axis of reform is the development of a graduated process of informed consent, not only based on the invasiveness of the neural intervention but also on the epistemic uncertainty of the long-term effects. For minimally invasive interventions, standard models of consent will suffice. But for interventions involving neural decoding, closed-loop stimulation, or pharmacologically mediated cognitive alteration, a multi-step framework of consent needs to be mandatory. This would involve a recursive model of consent, where subjects re-consent at specific neuropsychological points, with amended risk briefings based on real-time neural effect assessments.

Parallel to the reform in consent, a regime of neurodata protection needs to be legally codified. Neural data—especially that which is generated outside the clinical context—should be legally established as a protected cognitive biometric. The table below compares the current treatment of neurodata under relevant jurisdictions and proposes a legal paradigm for filing under U.S. BRAIN-compliant law:.

Finally, there is a need to rebalance the funding review process to make room for ethical metrics as weighted parameters in grant scoring. These metrics need to include neural justice effect (distributional equity), epistemic transparency (disclosure fidelity), and long-term autonomy

Table 1: Comparative Legal Treatment of Neural Data

Jurisdiction / Framework	**Neural Data Classification**	**Protection Level**
U.S. HIPAA	Not specified	Not protected unless generated in clinical care
EU GDPR	Ambiguously biometric	Partially protected if inferable to identity
Proposed BRAIN Act	Cognitive biometric	Fully protected; subject to consented usage, auditability, and revocability

preservation. Grant proposals need to be scored not only on technical novelty but also on how they minimize ethical entropy in cognitive ecosystems.

Taken together, these proposals are intended to render cognitive freedom a prime directive of federally funded neuroscience. The BRAIN Initiative not only must generate knowledge, it must encode democratic values into the design architecture of neural innovation. Anything less risks transforming the brain—the seat of individual personhood—into an unregulated frontier of commodified cognition and behavioral extraction.

References

1. National Institutes of Health. *BRAIN Initiative: Brain Research Through Advancing Innovative Neurotechnologies*. 2025. Available at: https://braininitiative.nih.gov/

2. NIH Advisory Committee to the Director BRAIN Initiative Working Group 2.0. *The BRAIN Initiative 2.0: From Cells to Circuits, Toward Cures*. National Institutes of Health, 2019.

3. Presidential Commission for the Study of Bioethical Issues. *Gray Matters: Topics at the Intersection of Neuroscience, Ethics, and Society*. 2014. Available at: https://bioethicsarchive.georgetown.edu/pcsbi/node/3543.html

4. National Science Advisory Board for Biosecurity (NSABB). *Guidance for Enhancing DURC Oversight*. National Institutes of Health, 2012.