European Commission

Research and Innovation

Ethics Advice Mechanism

European Group on Ethics in Science and New Technologies

Governing Neuro-AI

Towards an Infrastructure Approach

Independent Expert Report

Governing Neuro-AI: Towards an Infrastructure Approach
European Group on Ethics in Science and New Technologies

European Commission
Directorate-General for Research and Innovation
Unit 02 – Science for Policy, Advice and Ethics
Contact Mihalis Kritikos
Email EC-ETHICS-GROUP*@ec.europa.eu*
RTD-PUBLICATIONS@ec.europa.eu
European Commission
B-1049 Brussels

Manuscript completed in August 2026
First edition

Authors of the Statement (in alphabetical order): Nikola Biller-Andorno, Maria do Céu Patrão Neves (Vice-Chair), Migle Laukyte, Paweł Łuków, Fruzsina Molnár-Gábor, Thérèse Murphy, Barbara Prainsack (Chair), Nils-Eric Sahlin (Vice-Chair), Tamar Sharon, Jeroen van den Hoven, Takis Vidalis.

Print	ISBN 978-92-68-27432-3	doi:10.2777/0962282	KI-01-25-108-EN-C
PDF	ISBN 978-92-68-27431-6	doi:10.2777/8619693	KI-01-25-108-EN-N

Luxembourg: Publications Office of the European Union, 2026

EUROPEAN COMMISSION

ETHICS ADVICE MECHANISM

Governing Neuro-AI

Towards an Infrastructure Approach

European Group on Ethics
in Science and New Technologies

Statement No. 14

September 2026

2026 Directorate-General for Research and Innovation

Contents

1. Purpose and scope of this Statement

Recent reports and policy documents have identified important ethical and legal concerns surrounding neurotechnology and AI, including those relating to privacy, mental integrity, autonomy, consumer protection, and regulatory gaps. This Statement does not duplicate that work; instead, it foregrounds a strongly resonant but less developed and increasingly urgent issue: **the rise of advanced neuro-AI infrastructures.**

Advanced neuro-AI infrastructures are the integrated systems of hardware, software and associated practices through which neurodata are collected, processed, and reused to create, deploy, and manage AI-powered neurotechnology applications. These include the devices, data pipelines, models, cloud services, interfaces, business practices, and institutional arrangements that enable the continuous collection, interpretation, and modulation of human cognitive and behavioural processes through adaptive AI systems. Their potential is comparable to that of generative AI systems such as large language models (LLMs) that draw upon and shape the language, thoughts, emotions, and preferences of their users.

The development of neuro-AI infrastructures, the EGE argues, constitutes **a different governance challenge** from that traditionally posed by the processing and protection of individuals' data, not least because the infrastructures involved are not merely technical, but are embedded in political, economic, and moral choices. It thus requires us to move beyond a narrow focus on data governance and to recognise how the convergence of neurodata, neurotechnology, and systems with AI infrastructures requires a policy response that takes a systemic approach across a host of regulatory domains.

Neuro-AI systems – understood as systems that combine artificial intelligence with neurotechnologies to monitor, interpret, or interact with neural activity and cognitive processes – differ fundamentally from conventional digital systems because they can potentially infer latent mental states, adapt interactively, personalise influence, and shape cognition through continuous feedback loops. Many of the capabilities discussed in this Statement remain unevenly developed across different neurotechnologies and application domains. The governance concerns identified include the cumulative effects that may arise through continued advances in AI, data integration, and neurotechnology. Anticipatory governance should be proportionate to both present capabilities and plausible future developments.

Governance should therefore navigate not only data collection and processing, but also the broader societal implications of the creation and use of architectures, systems, and ecosystems capable of influencing cognition and behaviour. The latter we have termed the ***neurostack***, a coordinated ensemble of devices, data pipelines, Foundation Models, cloud services, interfaces, business practices, standards, audit and oversight mechanisms and institutional arrangements through which neurodata are collected, processed, reused, and operationalised (see also Key Terms). More specifically, we draw attention to the emergence of so-called Brain Foundation Models (BFMs) and related neurodata infrastructures that combine consumer neurotechnology, large-scale neurodata aggregation, transformer architectures, and downstream inferential systems. Unlike LLMs and other forms of generative AI, whose rapid deployment has led many societies and policymakers to conclude that opportunities for meaningful governance were missed, BFMs are still at a stage where their trajectories, institutional embedding, and societal uses can be actively shaped through democratic deliberation, public oversight, and anticipatory regulation. This may make it possible to subject BFMs and, more broadly, neuro-AI infrastructures to effective EU regulation and enforcement.

While ethical concerns regarding the convergence of neurotechnology and AI arise across medical, research, workplace, educational, security, and consumer settings, they are intensified by the promotion of commercial direct-to-consumer neurotechnologies and by the development and expansion of BFMs (in the form of so-called “Human (centric) Foundation Models” or HFMs) trained on large-scale neural, neuroimaging, behavioural, and contextual datasets. These developments point to an emerging ***political economy of neuro-AI infrastructures***, namely the political economy of integrated ecosystems of high performance computing, cloud storage, high speed and low latency communication channels and data pipelines, peripheral neurotech devices (i.e. those devices that interact with the peripheral nervous system rather than directly with the brain or spinal cord), and foundation models and inferencing tools, in which neurodata may become a durable resource for inference, prediction, classification, and consequential downstream uses.

The policy challenge is therefore **not limited to individual devices or downstream applications**. It also concerns a whole infrastructure of data pipelines, model architectures, institutional incentives, market structures, dual-use trajectories, and regulatory mechanisms through which neuro-AI systems are made possible.

In light of this challenge, this Statement adopts a broad risk-based approach that considers potential violations of rights and the undermining of moral values that these developments might entail. A wide range of possible harms ranging from physical harm to loss of agency and autonomy, and eventually undermining of dignity, need to be considered. Risk classification should depend on the intersection of several factors including inferential capability, adaptivity, autonomy, deployment context, and behavioural impact rather than focusing primarily on hardware invasiveness. Even relatively low-resolution non-invasive systems may become functionally invasive when combined with sufficiently powerful AI inferencing and adaptive behavioural optimisation.

At the same time, the Statement recognises the significant potential of neuro-AI to further scientific understanding and support medical developments, including improved access to health care, innovations in communication, and enhanced rehabilitation. It stresses that trustworthy infrastructures, scientifically valid claims, and robust governance aligned with the protection of fundamental rights, democratic accountability, and the public interest will all be required for such benefits to be achieved. Europe’s long-term strategic position will also depend on its ability to support competitive, trustworthy, and sovereign neuro-AI ecosystems.

The governance challenges discussed throughout this Statement arise at different stages in the operation of neuro-AI systems and should not be conflated. Some systems merely measure neural or neurophysiological activity; others generate increasingly sophisticated inferences about cognitive or emotional states; still others predict behaviour, adapt their interactions with users, or operate through closed-loop feedback that influences behaviour over time. These stages involve different technical capabilities, ethical concerns, and regulatory questions. The distinctive governance challenge addressed in this Statement concerns the ways in which such capabilities become integrated into reusable infrastructures whose cumulative effects may extend well beyond any individual device or application.

The authors thank everyone who has given input into this Statement, be it in discussions or in writing, notably:

Roberto Andorno, Jan Christoph Bublitz, Hervé Chneiweiss, Teresa De-Martino, Eugen Dolozal, Markus Frischhut, Kyriacos Hatzaras, Timo Istace, Sjors Ligthart, Antonia Mochan, Gloria Navarrete, Beth Parkin, Claudia Reinprecht, Christos Strydis, Yiannos Tolias, Guilherme Maia de Oliveira Wood, Katharine Wright and the SAEGE team.

1.1. Why is the EGE highlighting this issue right now?

This intervention by the EGE comes at a moment when neurotechnology and AI are converging. Neurotechnology increasingly relies on machine learning (ML) for signal interpretation, pattern recognition, personalisation, and adaptive functioning. At the same time, advances in AI and ML increasingly depend on access to large, heterogeneous streams of behavioural, biometric, and contextual data, gradually extending to a broad range of neurodata. The result is a growing co-evolution of neurotechnology and AI rather than two independent technological domains.

It is important that this convergence is accompanied, in parallel and in close coordination, by ethical assessment and policy development. Attention should not only be paid to whether a particular neurotechnology device is safe, effective, and privacy-preserving, but also to the wider questions of:

(1) how neural and behavioural data are collected, fused, aggregated, modelled, and reused across contexts (data governance);
(2) how such data contribute to reusable AI models including BFMs (model or infrastructure governance); and
(3) how downstream inferences may affect persons, groups, and institutions (societal governance).

In other words, neuro-AI, as a novel infrastructural technology, demands a systemic and structural policy approach.

1.2. What this EGE Statement contributes

Existing reports have addressed many important issues concerning neurotechnology and AI. These include privacy and integrity, freedom of thought, cognitive liberty, individual autonomy, human dignity, consumer protection, data protection, responsible innovation, and possible regulatory gaps (CCNE and CCNEN, 2026; EDPS, 2024; Geneva Academy, 2025; Global Privacy Assembly, 2024; OECD, 2019; OECD, 2025; STOA, 2024; UNESCO, 2025; WHO, 2026). This Statement builds on that existing work but shifts the emphasis from individual-level privacy and the protection of neurodata, and from isolated devices or applications, to the systemic, society-level impacts of neuro-AI infrastructures.

More specifically, this Statement distinguishes the model layer (BFMs/HFMs) from the data-and-pipeline layer (neurodata infrastructures) within the broader neuro-AI infrastructure. This focus makes visible questions that might otherwise remain dispersed across separate regulatory frameworks: for example, how neural and behavioural data are gathered and transformed into reusable model capacities; how commercial and institutional incentives shape development; how systems move across sectors; and how downstream inferences affect persons and groups beyond the original data subject. The Statement seeks to bring these concerns to the fore, reducing the risk that they only become the object of attention when it is already difficult and costly to effect change. This is particularly relevant for BFMs/HFMs and neurodata infrastructures, whose technical architectures, governance arrangements, and economic dependencies may become entrenched long before their societal consequences are fully understood.

This infrastructural focus helps to avoid two inadequate framings that dominate public and policy debates:

(1) It avoids reducing governance challenges to questions of individual privacy, informed consent, or data control, while fully preserving the importance of individual rights and their adequate protection. Although concerns about individual-level control over data remain essential, they are insufficient for understanding technologies that derive power not only from isolated acts of data collection but from the aggregation, integration, and continuous reuse of neurodata across platforms, devices, institutions, and contexts. The societal implications of BFMs/HFMs arise not merely from what happens to one person's data, but from the collective dynamics through which large-scale neurodatasets become the basis for predictive, adaptive, and potentially behaviour-shaping systems.

(2) An infrastructural perspective avoids treating commercial direct-to-consumer neurotechnology as a self-contained consumer-protection niche to be regulated on a consumer protection model. Devices marketed for wellness, productivity, gaming, education, or self-optimisation may simultaneously function as components of much larger neurodata infrastructures. In this context, consumer neurotechnology is not merely a collection of isolated products to be made available in a commercial context but part of a broader socio-technical assemblage through which neurodata are continuously extracted, standardised, aggregated, and rendered usable for future BFMs/HFMs and downstream inferential systems. Stated simply, products that appear relatively benign or low-risk when considered individually may therefore acquire very different ethical and political significance when viewed as elements within an expanding infrastructural ecosystem.

As noted at the beginning of our Statement, by 'infrastructure' we refer not only to technical systems in the narrow sense, but also to the wider arrangements that make certain forms of data extraction, modelling, coordination, and governance possible over time. These include hardware devices, software architectures, cloud infrastructures, datasets, standards, interfaces, business models, institutional practices, regulatory frameworks, and social norms. Whereas a 'system' may describe a bounded functional entity and an 'ecosystem' often emphasises networks of interacting actors and markets, the concept of infrastructure foregrounds the background conditions that become embedded, durable, and difficult-to-contest once established. Infrastructures shape social life precisely because they are 'invisible': they classify, standardise, and organise the world in ways that appear natural or inevitable and hard to change once widely adopted (Bowker and Star, 1999).

Viewing neuro-AI through an infrastructural lens therefore shifts attention from isolated devices or applications to the conditions under which entire forms of knowledge generation, prediction, and intervention become possible. It directs attention to questions of dependency, concentration of power, interoperability, path dependency, and asymmetries in access and control. It also highlights that infrastructures are not neutral technical backdrops but sites where political, economic, and normative choices are sedimented into durable arrangements.

A particular challenge of governing neuro-AI infrastructures lies in the combination of four features:

(i) the collection of neurophysiological data;
(ii) increasingly powerful inferential capabilities concerning cognitive and emotional states;
(iii) adaptive interaction with individual users; and
(iv) the potential for these capabilities to be reused across contexts through common models and infrastructures.

It is this combination, rather than any single technological component, that warrants distinct governance attention.

Importantly, adopting an infrastructural perspective also implies a shift in how policy and governance are designed. If the core risks arise not only from individual products but from cumulative infrastructures of data extraction, model development, and inferential capacity, then **governance cannot rely solely on downstream consumer protection, individual consent mechanisms, or case-by-case risk assessment**. Instead, policy must address upstream questions concerning data governance, control and access regimes, interoperability standards, public oversight, compute concentration, institutional accountability, and the societal purposes toward which neuro-AI infrastructures are directed. An infrastructural approach therefore calls for anticipatory and systemic forms of governance capable of shaping technological trajectories before economic and institutional lock-in make democratic intervention substantially more difficult.

1.3. Scope across sectors and contexts

Ethically relevant concerns arise equally across medical, research, educational, workplace, security, and military contexts. Commercial direct-to-consumer systems nevertheless deserve particular attention. They may combine large-scale data collection, consumer-facing claims, proprietary infrastructures, platform dependence, rapid iteration, closed loop feedback, poor traceability of neurodata in global hyperscale and cloud storage environments, and limited user understanding of downstream data uses. In such contexts, neural and behavioural data may be collected with consumer-facing promises of enhancements in areas such as sleep, concentration, productivity, mood management, gaming, accessibility or self-optimisation, while at the same time contributing to broader infrastructures of model training, classification, and prediction capabilities.

Commercialisation is therefore not merely a background or concomitant factor. It shapes which neuro-AI systems are developed, what kinds of data are collected, how claims are framed for consumers and investors, how quickly products are brought to market, and how neurodata are treated as strategic assets. As an example, the global market for commercial neurotech devices to enhance sleep, concentration, and other aspects of wellbeing is now estimated to rise to 44 billion dollars in the next 3 years. Start-up and venture capital activities are sharply increasing (Biospace, 2026). This does not mean that commercial actors are uniquely problematic, nor that applications in the public sector or in the wider domain of health care and medicine are free of risk. Rather, commercial incentives, market concentration, and platform dynamics should be considered as a constitutive part of the ethical and governance landscape.

1.4. Building on previous EGE work and EU values

This Statement builds on earlier EGE work on AI (EGE, 2018) and on democracy and digital technologies, including concerns about platform power, algorithmic influence, asymmetries of knowledge, and democratic accountability (EGE, 2023). Neuro-AI infrastructures raise similar questions in a particularly sensitive domain, because they may support inferences about cognitive, affective, behavioural, or health-related states. Once infrastructures, standards, data pipelines, and foundation models become widely adopted, they may generate forms of technical, economic, and institutional lock-in that are difficult to reverse (Bowker and Star, 1999).

This Statement is grounded in existing commitments of the EU and its Member States to human dignity, autonomy, privacy, non-discrimination, freedom of thought, bodily and mental integrity, scientific responsibility, and democratic accountability. It does not presuppose the need for a separate catalogue of neuro-specific rights. **Its central concern is how existing rights and ethical principles should be operationalised in relation to emerging neuro-AI infrastructures.** These developments also raise questions of **European technological sovereignty, public-interest infrastructures, and democratic accountability.**

2. From neurotechnology to neuro-AI infrastructures

Individual-level neurotechnologies are evolving into a data-generating layer inside wider data–model–platform ecosystems that enable continuous inference and reuse across contexts. This shift is driven by multiple, concurrent developments that together create a strong path from neurotech to neuro-AI infrastructures:

a. The move from episodic measurement to persistent, reusable inference across contexts (section 2.1)
b. The rise of BFMs as a new infrastructural layer trained on large, heterogeneous neuro-related datasets (section 2.2)
c. The expansion of consumer-facing neurotech and wearables that function as pipelines for large-scale data aggregation and model training (sections 2.3 and 2.4)

d. An infrastructural re-framing that highlights upstream data governance, interoperability, platform power, and lock-in beyond single devices or apps (section 3.1)
e. Cloud-based processing and aggregation that treat brain and behavioural data as assets for training reusable models at scale (section 3.2)
f. Cross-domain “dual-use” transfer and reconfiguration of models and systems across medical, workplace, educational, security, and consumer settings (section 3.5)

Taken together, these dynamics call for governance that addresses not only individual- and device-level safety and efficacy, but also the upstream conditions of data collection, model training, platform integration, and downstream institutional use across sectors.

2.1. From episodic measurement to persistent inference

Earlier forms of neurotechnology were often device-centred and context-specific. Clinical electroencephalogram (EEG), deep brain stimulation, assistive brain-computer interfaces, and neurorehabilitation technologies have typically been developed within research, medical or therapeutic settings, with relatively well-defined purposes and institutional oversight and limited to one type of neurodata.

This is now changing, with neurotechnologies increasingly serving as data-generating nodes within broader data–model–platform ecosystems, where neural signals are aggregated, fused, and reused to build BFMs and other downstream inferential tools. Neurodata may thus no longer be collected only for a bounded clinical or research purpose. They may become part of longitudinal datasets, training pipelines, and reusable models that support future inferences, classifications, or interventions in other contexts.

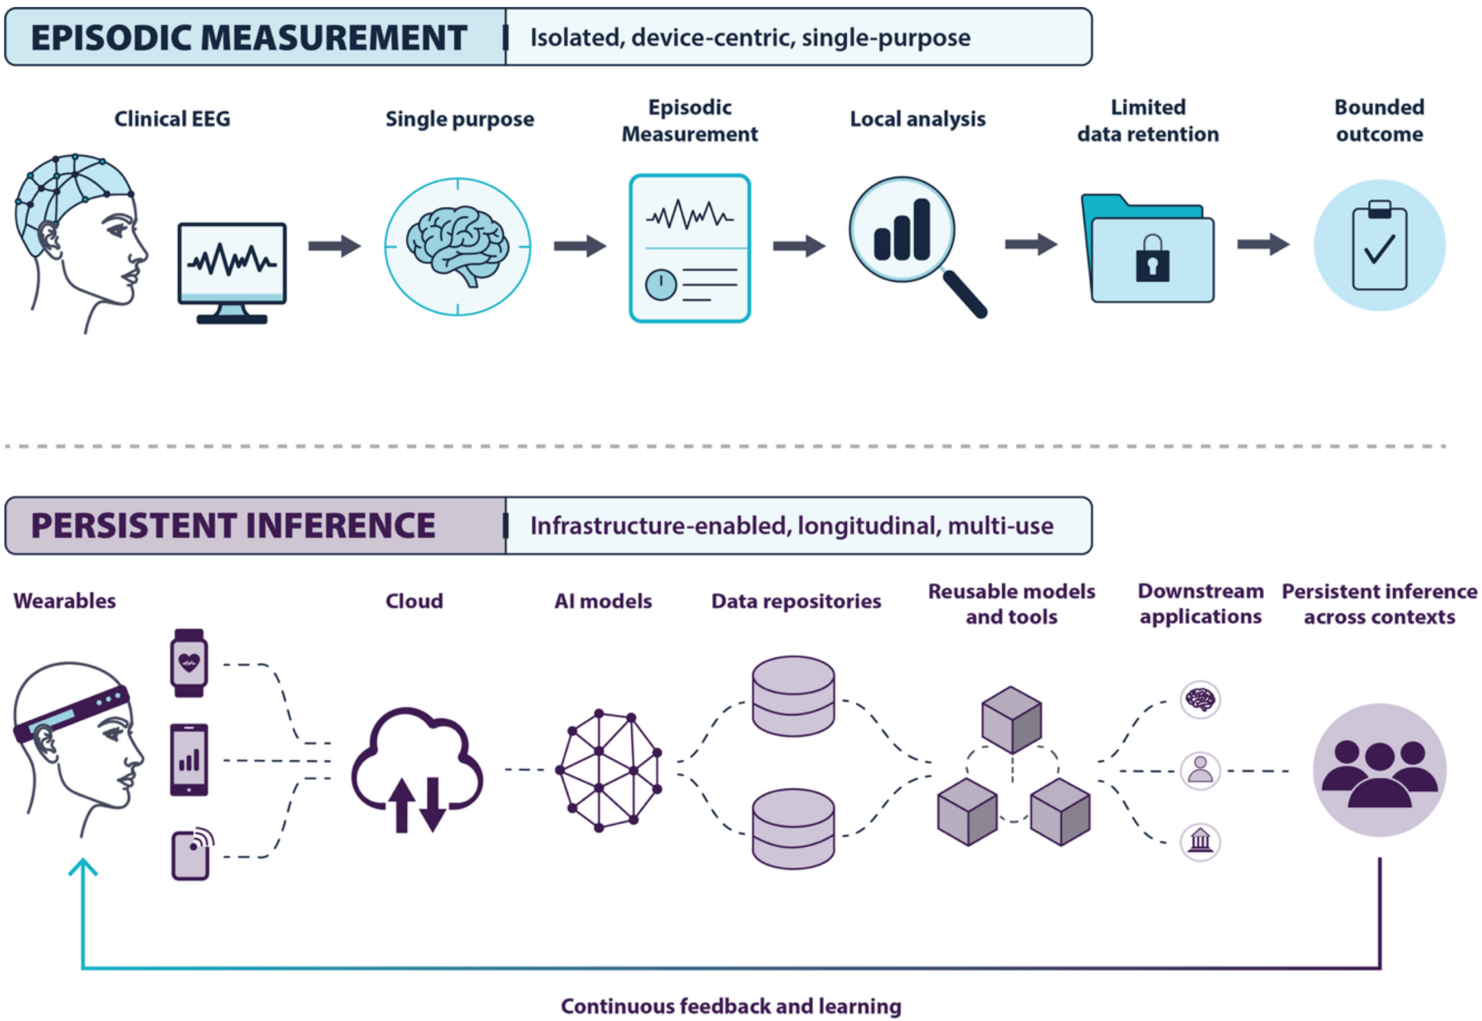

Figure 1: From episodic measurement to continuous inference. Differences between the two forms of neurotechnologies. (Created by SAEGE)

2.2. Brain Foundation Models as an emerging infrastructural layer

BFMs are among the most prominent manifestations of this development. BFMs are AI models pre-trained on large, heterogeneous, and often un-labelled neuroimaging datasets spanning electrophysiological, functional, and structural modalities, including EEG, MEG (magnetoencephalography), fMRI (functional magnetic resonance imaging), MRI, and PET (positron emission tomography). Although still at an early stage of development, BFMs seek to learn transferable representations that can then be adapted to multiple downstream tasks. They could mean a shift in neuroscience from highly specialised, task-specific tools towards general-purpose, reusable AI models with the potential to support a range of downstream applications in clinical diagnosis, neurorehabilitation, and cognitive neuroscience research (see Box 1 below). At the same time, BFMs may also support highly specialised uses, including precision neuroimaging and patient-specific analyses (Dibble et al, 2026). Their significance therefore lies not only in their technical novelty, but in their potential to convert distributed neurodatasets into reusable and integrated infrastructures for classification, prediction, and downstream decision-support (Hanley et al., 2026; Kwon and Shin, 2026). Analogous to the popularity of benchmarks and criteria for assessing the quality and performance of LLMs and coding models, there are now competitions for BFMs, and Github curates a list of available BFMs (Github, undated).

BOX 1: Beacon

MIT News reports on a recent business start-up named Beacon. Beacon partners with pharmaceutical companies to accelerate its path to patients. The company's FDA-cleared medical device has already been used in over 40 clinical trials across the globe as part of studies aimed at treating conditions including major depressive disorder, schizophrenia, narcolepsy, idiopathic hypersomnia, Alzheimer's disease, and Parkinson's disease. With each deployment, Beacon learns more about how the brain works – insights it is using to create a "foundation model" of the brain (Winn, 2026).

At the time of the writing of this Statement, the development of BFMs faces significant obstacles, including limited data volumes, substantial inter-subject variability, weak standardisation, and uncertainty regarding scaling laws - the empirical, roughly power-law pattern where a model's performance improves predictably as model size, data, and compute increase - though with diminishing returns for each additional unit. At the same time, a recent pre-publication on training-data governance for BFMs argues that they already raise distinctive governance questions because of their scope to repurpose neurodata historically collected in clinical and research settings through practices such as large-scale reuse, cross-context combination, and open-ended downstream application (Hanley et al., 2026).

As with foundation models in other domains, the governance challenge of BFMs is not confined to a single domain; it spans training-data provenance and the legal basis on which data have been obtained, collection conditions, model documentation, access and release modalities, downstream adaptation, and the distribution of benefits and risks. This is why BFMs require attention not only as research artefacts or clinical decision support tools, but as potential model infrastructures that operationalise neural and behavioural data across sectors.

2.3. Consumer wearables as an example of infrastructure consolidation

BFMs are a fundamental layer in what could be called an emerging neurostack. A neural foundation model can communicate with peripheral devices such as AI-enabled smart glasses (which can in turn communicate with smart watches) or sleep headbands or even caps. This provides a useful example for understanding the shift to which we want to draw attention. It illustrates how a consumer-facing wearable can function simultaneously as a sensor, interface, cloud-processing node, and model-training pipeline. The immediate service may be sleep improvement, communication, translation, navigation, memory support or accessibility. The broader infrastructure, however, may involve continuous multimodal data collection, AI interpretation, human or automated review, model improvement, and integration into platform ecosystems as illustrated in Box 3 below Anzolin and Lo Nostro, 2025; Lepapa et al., 2026; Reuters, 2025).

BOX 2: The example of the sleep/wellness data pipeline

A consumer uses a sleep-enhancement EEG headband. The immediate service is sleep feedback. Aggregated and fused data may also support model training, product development, risk prediction or downstream mental-health inferences. The ethical issue is not only whether the user consented to sleep feedback, but whether data collected in this context may become part of a reusable infrastructure for future inferences about that user or others (Hanley et al., 2026; Horikawa et al., 2013). The sleep neurotech applications are seen as a very promising and lucrative market as many US Americans use sleep medication and many Americans are heavily dependent on them (Reuben et al., 2023).

This example of wearables is especially relevant because wearable AI systems may converge with more active or agentive neurotechnological interfaces. For example, Meta's acquisition of CTRL-Labs, a company developing wrist-based electromyography systems intended to translate neural signals into device control, suggests a possible trajectory in which smart glasses, wearable neural interfaces and AI systems might become integrated into unified consumer interfaces (Constine, 2019). Such systems may increasingly resemble non-invasive brain-computer interfaces, even when they are not marketed as medical or neurotechnological devices.

Consumer-facing neuro-adaptive technologies may become increasingly important. Examples include AI-enabled wearables, smart glasses, neural interfaces, wellness applications, adaptive tutoring systems, and personalised cognitive-assistance technologies. These systems may continuously collect data, adapt their behaviour to individual users, and optimise interactions in ways that blur the distinction between assistance, personalisation, and behavioural influence.

The principal concern is not necessarily the ability to "read minds", but rather the possibility of continuous sensing combined with continuous adaptation under commercial incentives. Such systems may create new forms of personalised persuasion, behavioural optimisation, engagement maximisation, or attention shaping, thereby raising questions concerning autonomy, consumer protection, and democratic accountability.

The point of this example is not to claim that current smart glasses constitute neuro-AI systems. It is to show how consumer devices can evolve into infrastructures of sensing and inference and ultimately behavioural interaction. The governance challenges become more acute when analogous infrastructures incorporate neural or neurophysiological data.

2.4. Looking beyond BFMs

Several other emerging technologies and infrastructures may raise equally significant, and in some cases even more immediate, ethical, legal, and societal challenges. The governance of neuro-AI therefore cannot be limited to BFMs but must also consider the wider ecosystem of technologies through which neural and neurophysiological data are collected, interpreted, integrated, and acted upon. A number of these are set out below, with an indication of the specific governance challenges that they pose.

2.4.1. Adaptive closed-loop neurotechnology

One important development is the emergence of adaptive closed-loop neurotechnologies. Traditional neurotechnology often functions as a sensing or measurement device: neural activity is recorded, interpreted, and used to support diagnosis, therapy, or research. Closed-loop systems, by contrast, continuously adapt their behaviour in response to user data. Neural signals may be processed by AI systems that dynamically modify stimulation parameters, feedback mechanisms, user interfaces, or behavioural interventions in real time.

Such systems, therefore, do not merely observe neural activity but are likely to actively influence attention, behaviour, emotion, motivation, or decision-making processes. The governance challenge extends beyond data protection and safety to questions of autonomy, manipulation, transparency, accountability, and meaningful human control. As adaptive systems become more sophisticated, the distinction between observing cognition and shaping cognition may become increasingly difficult to maintain.

2.4.2. Neuro-AI digital twins

A second development concerns the emergence of neuro-AI digital twins (see Box 3). These systems combine neurodata with behavioural, physiological, clinical, and contextual information to construct computational representations of individual persons. Such models may eventually support treatment planning, prognosis, neurostimulation optimisation, rehabilitation, or personalised therapeutic interventions.

BOX 3: Example of neuro-AI digital twin development

Early in 2026, Meta AI released its Tribe v2 Neural Foundational Model that can be considered as a paradigm case of neuro-Ai digital Twin development. This model serves as a digital twin of the human brain and a general-purpose encoding model of whole-brain activity across cognitive domains, embedding video, text and audio inputs, integrating them via a Transformer component into fMRI voxel activation predictions. According to the claims in the introductory text, it is a “foundation model capable of predicting human brain activity in a variety of naturalistic and experimental conditions. Leveraging a unified dataset of over 1,000 hours of fMRI across 720 subjects, we demonstrate that our model accurately predicts high-resolution brain responses for novel stimuli, tasks and subjects, superseding traditional linear encoding models, delivering several-fold improvements in accuracy”. The models train themselves on the input provided by masking, ablation, prediction and error minimisation (Kang et al, 2025). They no longer require human experts to provide the ‘ground truth’ for the classification of each sequence of images or recordings or raw neurodata, thus speeding up the development of the models by orders of magnitude (Meta, 2026).

The governance significance of digital twins lies in the possibility that they may become more informative than the underlying data from which they are derived. Questions arise concerning control, consent, model validity, liability, transparency, and the permissible uses of individualised representations of persons. Governance frameworks may therefore need to address not only the collection of neurodata but also the creation and use of computational representations derived from those data.

2.4.3. Multimodal human foundation models

A third development is the emergence of multimodal human (centric) foundation models (HFMs) that combine neurodata with behavioural, physiological, biometric, linguistic, audiovisual, and contextual information. In such systems, neurodata may no longer constitute a distinct category of information but rather one component within a much broader representation of human behaviour and cognition that may serve as a versatile baseline that can be fine-tuned for specialised applications such as digital avatars, robotics, medical research, and social intelligence (Tang et al., 2025).

This development raises questions that extend beyond neural privacy. Governance concerns increasingly relate to inferential privacy, profiling, re-identification, proxy inference, and the construction of highly detailed computational representations of individuals. The object of governance may gradually shift from neurodata themselves to the integrated models that represent persons across multiple dimensions of behaviour and experience.

2.4.4. Cognitive-state inference systems

Emerging technologies increasingly seek to infer cognitive and emotional states such as attention, fatigue, workload, stress, emotion, engagement, or intention. Such systems may rely on combinations of EEG, eye tracking, wearable sensors, behavioural monitoring, speech analysis, or other physiological signals.

Besides asking whether such inferences are accurate, we also need to ask whether it is ethical for any type of actor to use these inferences, including consideration of by whom and under which conditions. Specifically, questions arise concerning fairness, discrimination, transparency, due process, and the legitimacy of decisions based on probabilistic assessments of cognitive states. Such concerns may become particularly relevant in workplaces, educational settings, military environments, and other contexts that are consequential and characterised by asymmetries of power.

In summary, taken together, these developments suggest that BFMs constitute one layer within a broader emerging neuro-AI infrastructure (which we call the neurostack) that includes adaptive systems, digital twins, multimodal foundation models, cognitive-state inference technologies, platform infrastructures, and consumer-facing applications. The most significant governance challenges arise not from any single technology, but from the interactions between these layers and the ways in which they transform neural and behavioural data into reusable, scalable, and socially consequential capabilities.

3. The political economy of neuro-AI infrastructure

As discussed earlier, neurotechnologies were initially developed and governed primarily as individual devices in specific clinical or research contexts, with attention paid to factors such as safety, accuracy, invasiveness, medical purpose, and user protection. Today, these devices increasingly serve as entry points into neuro-AI infrastructures in which neural signals are aggregated, fused, and reused for BFMs, HFMs, and other downstream inferential tools.

In this infrastructural mode, consumer neurotech and adjacent wearables expand data collection into everyday settings, feeding large-scale, heterogeneous datasets for model pretraining and reuse. The infrastructural perspective helps to make visible how these data flows connect devices to cloud services, standards, interfaces, and business models, creating background conditions that shape knowledge production, dependency, concentration of power and path dependency over time. As a result, neurotechnology must be understood as a data-generating component within broader socio-technical systems – where governance focuses not only on product risk, but also on upstream data governance, model and platform layers, and societal impacts of downstream inferences.

3.1. Neurodata as a strategic resource

The emergence of neuro-AI infrastructures should be understood not only as a technological development, but also as a political-economic one. Neurodata may become valuable by providing a specific service to a user and also by training models, improving products, developing proprietary datasets, enabling prediction, supporting classification, and creating market advantage.[1]

The political economy of these systems extends well beyond firms' own profit-seeking behaviour to incorporate issues including the organisation of data access, compute capacity, intellectual property requirements, cloud infrastructures, platform distribution, investor expectations, public procurement rules, and regulatory classification. All of these – separately and jointly – influence which applications are developed, which populations become data sources and hence set standards and patterns of representativeness, which benefits are prioritised, and which risks remain externalised.

This systemic or socio-technical framing is important because it leads us to recognise that ethical risks do not arise only from misuse or bad actors. They may also arise from well-known institutional incentives: the incentive to collect more data than is needed for the immediate service; to retain data for future model development; to make strong claims in competitive markets; to avoid medical classification and associated liabilities; to rely on opaque infrastructures and data storage; or to repurpose data across contexts.

Governance therefore needs to address the conditions under which neuro-AI systems are built (from the beginning), rather than focusing solely on remedying harms after deployment. A 2024 analysis of the data policies of 30 consumer neurotech companies by the Neurorights Foundation, a non-profit organisation in New York City, for example, showed that nearly all had complete control over the data users provided. This means that most firms could use the information as they please, including commercial exploitation (Genser et al., 2024).

3.2. Neuro-AI infrastructures as data pipelines

Consumer-facing neurotechnologies typically operate on a notice-and-consent model that is formally valid yet substantively hollow: users click "agree" without any realistic grasp of how their neural data will be used, trained on, repurposed, or transferred. Individual consent may thus not be sufficient to justify infrastructural risks related to emerging data ecosystems as they affect the individuals concerned, let alone to take into account the broader societal implications of converting dispersed neurodata into reusable model infrastructures.

UNESCO and BCC studies of the neurotech market show a 700% increase between 2014 and 2021 and estimate a global investment level in neurotech of circa $24 billion by 2027 (UNESCO, n.d.; BCC Research, 2023). A recent study by the Centre for Future Generations analysed 271 neurotechnology companies worldwide, systematically distinguishing between medical and consumer applications, to provide an overview of this emerging sector (Centre for Future Generations, 2026). They reported a rapidly evolving consumer landscape with lean start-ups, fast market entry, and integration into mainstream wearables. There are dozens of new brands in the wellness and fitness business straddling a regulatory grey-zone, with claims bordering

[1] See, for example the forthcoming programme of the US National Institutes of Health that "seeks applications for an interconnected set of knowledge base projects including a data commons that will federate currently established BRAIN data repositories into a unified framework" (National Institutes of Health, 2026).

on the medical such as “reduce stress in minutes, forget fatigue”. These companies are pursuing revenue streams via the acquisition of consumer neurodata and via AI models that are (pre) trained on these data, and that will later provide the basis for a variety of commercial products and services.

3.3. Market concentration, platform power, and dependency

BFMs, HFMs, and related inferential tools are likely to concentrate strategic advantage in actors with privileged access to large-scale neurodata, computational resources, cloud infrastructures, model architectures, distribution channels, and downstream markets. These capabilities are distributed highly unevenly and tend to further entrench the position of actors already embedded within powerful technology ecosystems or able to aggregate and integrate data across consumer, research, healthcare, workplace, and platform environments. As a result, the development of BFMs and related systems may intensify existing asymmetries in economic power, epistemic authority, and infrastructural control.

This raises concerns about market concentration, the deepening of unreliable and unsafe dependencies, regulatory arbitrage, public-interest capacity, and digital sovereignty. One of the largest concerns is the rapid normalisation and commercialisation of large-scale neurodata collection before adequate governance structures, safeguards, and accountability mechanisms are established. Once such infrastructures become economically entrenched and socially normalised, appropriate governance may become substantially more difficult (Boonstra, 2025; Ienca et al., 2022; UN Human Rights Council Advisory Committee, 2024).

Recent years have clearly demonstrated how LLMs and foundation models in general (introduced in the market as AI Chatbots, such as ChatGPT, Perplexity, Deep Seek, Gemini, Claude, and Microsoft Co-Pilot) have been rapidly taken up (an estimated 1.5 billion people frequently use them) and integrated with widely used and accepted applications. It is also evident by now that they bring new risks of harm and opportunities for manipulation, exploitation, deception, misinformation, dependencies, and bias. These statistical Deep Learning next word prediction models have interesting applications, but they have also misinformed and misguided their users. They have given rise to a range of mental health issues, safety, and security problems, and brought great power and market dominance to their owners (Olsen et al., 2026). Transposed to the neuro-domain, many of the ethical issues associated with LLMs will surface here as well.

If neuro-AI infrastructures are shaped primarily by proprietary and opaquely accumulated datasets, non-EU cloud infrastructures, closed models, and platform logics, public authorities may later have limited capacity to assess, contest, or redirect the systems on which future health, workplace, educational, or security applications may depend.

These concerns do not imply that public-sector development is inherently preferable or that private innovation should be constrained as such. Rather, experience from adjacent digital domains – including platform economies, large-scale AI development, and cloud infrastructures – suggests that early advantages in access to data, compute capacity, standards-setting, and distribution channels can generate durable forms of market concentration and infrastructural dependency. In the emerging neuro-AI domain, governance must therefore attend to structural asymmetries: who is able to collect and aggregate neurodata at scale, who develops and controls foundational models, who defines technical and scientific standards, who has the authority to validate performance and safety claims, who captures the benefits of downstream applications, and who is exposed to the social, economic, and political risks associated with their deployment.

As we explored in our 2023 Opinion, Democracy in the digital age (EGE, 2023), society is subject to new forms of "surveillance capitalism" (Zuboff, 2019) in which vast quantities of data relating to our lives are accumulated, collated, and repurposed for commercial ends. We may now be transitioning towards forms of "neuro-adaptive capitalism" in which systems continuously optimise behavioural influence through neural inferencing and adaptive interaction. While surveillance capitalism concerns the accumulation of data relating to people's waking lives, it has been argued that these encroaching forms of neuro-capitalism additionally target sleeping subjects and concern sleeping and dreaming (Kane, 2025; see also Box 4).

To date, governance structures have tended to focus primarily on the collection, processing, and transfer of data, making them less well suited to addressing the broader interests implicated by such data use, including the impacts of continuous cognitive influence. This mismatch between data governance and the governance of persistent cognitive influence may ultimately become one of the defining governance challenges of neuro-AI systems. There are, however, some promising developments that provide a foundation on which to build. The AI Act already contains prohibitions on certain manipulative and exploitative uses of AI that, as we explore further in Section 5.2, may apply to some forms of neuro-AI infrastructures, although it does not explicitly address the more challenging issue of cross-contextual re-use. Likewise, the Digital Services Act (DSA) offers a useful model through its requirements for systemic risk assessment by very large online platforms, including risks relating to fundamental rights, civic discourse and electoral processes, manipulation of services, and transparency around recommender systems.

BOX 4: Targeted dream incubation

Recent work on so-called 'targeted dream incubation' and its commercial use proves that legislative developments regarding AI of the EU were highly prescient. A group of the top sleep researchers in the world published a letter of concern about advertising projects utilising these techniques, drawing attention to attempts to steer dreams of sleeping subjects for marketing purposes (Haar Horowitz et al., 2021). However, the fact that trauma-related nightmares may also be treated clinically with the same techniques points again to the importance of a balanced and proportionate governance posture that encourages these types of beneficial uses of neuro-AI in the public interest (Youngren et al., 2025; see also Marlan, 2023).

3.4. Complementing individual consent: public-interest and societal implications

An inherent feature of neuro-AI infrastructures is the way that they regularly generate effects that extend well beyond the individual data subject. They may, for example, shape classificatory categories (such as “prone to aggression”, “depression”, or “impulsivity”), contribute to the training of reusable models, enable the creation of group-level profiles, influence institutional decision-making, and alter the conditions under which other people are assessed and treated. For this reason, even valid individual consent to the use of particular datasets cannot by itself justify large-scale neurodata aggregation and the creation of BFMs/HFMs.

Governance frameworks therefore should not only maintain respect for individual rights, but also embrace a broader perspective that considers the social, collective, and infrastructural consequences of data processing, including impacts on groups, institutions, and shared public goods (Prainsack et al., 2022). The challenge is not to weaken individual control and protection, but to complement the existing DSA systemic risk framework (democracy, public discourse, societal harms) and the AI Act’s high-risk and systemic risk logic with mechanisms that recognise the relational and societal dimensions of neurodata infrastructures and ensure that their benefits, burdens, and risks are governed in ways that serve the public interest (Prainsack, 2026).

3.5. Dual use and cross-domain transfer

Cross-contextual use is a structural feature of neuro-AI infrastructures. The same datasets, models, and inferential systems may move across medical, consumer, workplace, educational, security and, in some cases, military contexts. Systems developed for health, accessibility, or wellness may later be repurposed for monitoring, classification, selection, operational performance, or security applications. This cross-sectoral use includes, but is not limited to, what EU law characterises as “dual use”, namely technologies that may serve both civilian and military purposes. Because of this, governance cannot rely only on the declared purpose of a single device or application (Hanley et al., 2026; Mantellassi and Madziwa, 2025; Stinchfield, 2023). It must address downstream adaptation, transferability, access conditions, procurement contexts, and institutional reuse.

The possibility that some neuro-AI systems acquire cross-contextual uses should not be understood as a reason to impede societally beneficial innovation. It underscores the need for governance frameworks that anticipate downstream adaptation and changing contexts of use. Such frameworks should combine appropriate export controls where military applications arise with broader requirements for documentation, risk assessment, access governance, post-deployment monitoring, and accountability throughout the lifecycle of neuro-AI systems.

Special attention is also required for the rapid uptake and development of neurotechnology designed specifically for military purposes. Historically, one of the major public funders of neurotechnology has been the US military through the Defense Advanced Research Projects Agency (DARPA), whose investments in brain-computer interfaces (BCIs) have generated technological advances with downstream applications in civilian and consumer devices. Military applications continue to expand. Examples include DARPA's Next-Generation Nonsurgical Neurotechnology (N3) programme, which seeks to develop high-performance, bidirectional brain-machine interfaces for able-bodied service members; sixth-generation fighter aircraft incorporating non-invasive BCIs into pilot helmets; neural interfaces integrated into military exoskeletons; and experimental systems enabling the control of autonomous platforms through BCIs (DARPA, n.d.; Greenbaum, 2016). China is likewise reported to be investing in AI-enabled military neurotechnologies (Finlay and Hooten, 2026). As BFMs and related neuro-AI technologies mature, they may confer significant strategic and operational advantages in conventional warfare, cyber operations, hybrid conflict, and cognitive warfare. At the same time, these developments raise important legal and ethical questions, including challenges for compliance with International Humanitarian Law (Gielas, 2025; see also Mantellassi and Madziwa, 2025).

China has officially made BCIs – both civil and military – into one of the pillars in its 15-year future industries programme (Finlay and Hooten, 2026). By elevating BCIs to this status, China turns the human brain into a domain of strategic global competition. Technologies that connect human neural activity directly with machines, weapons systems, or networks have become respectable scientific and R&D aims: capabilities to be developed, integrated, and scaled.

This geopolitical dimension also raises important questions about the cybersecurity of neuro infrastructure and neurodata, especially given the recent development of very powerful foundation models that may lead to geo-political asymmetries in the neuro-domain (Dey et al., 2026).

As a result, EU investment in neurotechnology through defence, security, and competitiveness funding instruments should be guided by a specific ethical and fundamental rights framework that takes account of both civilian and dual-use applications. Such a framework should require prospective assessment of downstream adaptation, cross-contextual deployment, and potential impacts on cognitive autonomy, while ensuring that innovation serving legitimate public purposes can continue to be developed responsibly.

4. Ethical and fundamental rights concerns

In this section we identify ethical and fundamental rights concerns including the wider societal implications of neurodata use across multiple contexts, elaborating on the governance challenges identified in section 3. Many of these concerns also arise across neurotechnology contexts more broadly. However, the convergence of large-scale data aggregation, foundation-model architectures, platform infrastructures, and commercial incentives may intensify these concerns or create new governance challenges.

4.1. Neurodata and wider societal implications

As we have reiterated throughout this Statement, neurodata may contribute not only to decisions concerning the primary data subject (i.e. the person that the neurodata were initially colleected from), but also to the development of reusable models and inference systems. Thus neurodata and neurodata-derived inferences require heightened protection where they may reveal or support wider inferences about mental, cognitive, conative, affective, behavioural, or health-related states. Their ethical significance does not, therefore, arise only from individual sensitivity. It also emerges through aggregation, comparison, combination with other data, and reuse in model development (CCNE and CCNEN, 2026; EDPS, 2024; Rainey et al., 2020; Rainey 2023; Susser and Cabrera, 2024).

This creates risks for individuals, groups, and populations. Neurodata from one set of users may be used to train models that classify or affect other individuals who did not contribute to the dataset on which the analysis and model development are based. This is why governance needs to go beyond individual privacy, the validity of consent or the level of control exercised by the data subjects. Emerging risks cannot be managed by reference solely to individual rights, including the protection of individual autonomy. They extend to wider societal risks, including those related to collective profiling, group-level bias, social sorting, and institutional decision-making based on neural or behavioural proxies.

Governance should therefore consider not only the proportionality of data collection, but also the proportionality of inference. The permissibility of increasingly powerful neurodata-derived inferences should depend on their scientific validity, intended purpose, context of use, foreseeable consequences, and impacts on fundamental rights. This principle of inferential proportionality complements, rather than replaces, established principles such as data minimisation, purpose limitation, and proportionality in EU law. It is especially relevant where relatively limited or low-resolution neurodata are combined with behavioural, biometric, contextual, or health-related data to generate consequential inferences about persons or groups.

4.2. BFMs/HFMs and cross-context reuse

BFMs/HFMs raise specific concerns because they may be trained on data collected across different contexts and then adapted for downstream purposes not anticipated at the point of data collection. Hanley et al. identify privacy, consent, bias, benefit sharing, and governance as central issues for training-data governance in BFMs (Hanley et al., 2026). These concerns include training-data provenance, lawful basis, benefit sharing, third-country processing, model release, and downstream adaptation.

As we have seen, data collected for meditation, gaming, or sleep optimisation may, for example, later contribute to a model adapted for clinical, educational, commercial, or security contexts (see section 3.2). The ethical and legal issue at stake is not only whether the original user agreed to a specific service or to the processing for a different purpose, but whether data collected in one context may become part of a reusable infrastructure for future inferences about other individuals in other contexts (Guenther et al., 2024; Hanley et al., 2026; Pan et al., 2017). The ethical and legal concern therefore extends beyond issues related to purpose limitation and the informed consent of the original data subject for re-use in other contexts. It concerns the transformation of context-specific data into reusable infrastructures capable of generating inferences across multiple domains. **Reuse thus changes the object of governance.**

4.3. Epistemic overreach, hubris and injustice

A central concern in the context of neurotechnology and predictive power is epistemic overreach: the risk that neural or behavioural inferences are treated as more authoritative than persons' own accounts of themselves, their deliberative choices, or contextual self-understanding. This concern is closely related to epistemic hubris: overconfidence in what neurodata and AI systems can legitimately reveal about persons, intentions, preferences, or capacities (Pongiglione, 2025). No amount of neuro-scientific information about an individual, even if highly accurate, is entirely "complete" in the sense that it represents that individual in a comprehensive and meaningful way. However informative, it cannot be legitimately considered to represent a specific person or capture what it is like for a particular person to be in a particular mental state. Accordingly, the sort of claim that one could "know" a person through neurodata (combined with other data) does not hold up other than in a metaphorical or diminished sense. And the same applies to claims that "people are their brains" (Musk, 2025), or that the natural language they use to talk about themselves and their experiences is just a second-hand encoding of the "real thing", namely brain states.

This bypassing of the person's own perspective gives rise to a further concern of epistemic injustice. Neuro-AI systems may displace or downgrade persons' own accounts of their experiences, intentions, preferences, or capacities when neural or behavioural inferences are treated as more objective or authoritative. This risk is especially serious where such inferences are used in consequential contexts, including healthcare, education, employment, insurance, law enforcement, or security. Even where systems are statistically useful, their outputs should not be allowed to substitute for situated judgement and contextual interpretation.

If a neural or behavioural inference about stress, attention, impulsivity, or credibility were to be treated as more reliable than a person's own account in, for instance, an employment, medical, insurance, educational, or legal context, the ethical problem is not merely possible inaccuracy, but the institutional displacement of a person's own account by model-generated inferences, thereby objectifying and dehumanising the person. Furthermore, should model-generated inferences become the basis of interventions in the form of neural stimulation and modulation, this would raise questions about who could legitimately exercise this control, in what contexts, and within what kinds of parameters.

4.4. Manipulation, influence and consequential uses

Neuro-AI infrastructures may support forms of prediction, targeting, persuasion, or intervention that affect autonomy, dignity, freedom of thought, bodily integrity, and mental integrity. The risk associated with these uses is not limited to intentional manipulation. It also includes the normalisation of systems that classify, rank, nudge, or influence persons through uncertain or weakly validated inferences (Alegre, 2021; Bublitz, 2014; OSCE ODIHR, 2023; Yuste et al., 2017). Neuro-AI systems may also influence behaviour through optimisation logics that seek to maximise engagement, learning, compliance, productivity, or health outcomes. Even where such objectives are beneficial, the resulting interventions may shape choices, opportunities, and patterns of behaviour in ways that raise questions of autonomy, accountability, and legitimate influence (Timms, 2025).

Particular concern arises in contexts of undue dependency, vulnerability, or structural power asymmetry: for example with respect to impacts on children, patients, persons with mental health conditions, workers, students, persons subject to law enforcement, or persons dependent on access to insurance or public services. People in certain contexts could feel they have little choice but to use certain neurotech applications or could be pressured to use them.

Work on immersive and emotionally salient digital environments further suggests that convergence between neurotechnology, virtual reality, and AI may create distinctive vulnerabilities in contexts involving dependency, persuasion, or altered perception (Kellmeyer et al., 2019). Cognitive liberty is a key consideration as it can be restricted by these types of models. However, it is just as important that these models do not influence our **emotions and moods**, in an undesirable or involuntary way, e.g. by making it difficult to mentally represent and put our feelings into words. These types of models should not be allowed to influence our emotional liberty.

4.5. Public-interest governance and digital sovereignty

Because neuro-AI infrastructures may shape societal categories, institutional practices, and collective conditions, they require democratic accountability, public-interest research capacity, and robust European governance. They are a part of a common epistemic and cognitive infrastructure that should be under governance, oversight, and public interest stewardship.

In the EU context, this is an important dimension of strategic autonomy. If neuro-AI infrastructures are primarily shaped by commercial or state entities outside of the EU, based on their proprietary datasets and market concentration, European values may be implemented only after key infrastructures have already been built elsewhere (Centre for Future Generations, 2026; Helbing and Ienca, 2024), possibly limiting their appropriate realisation.

Governance that regulates neuro-AI in the public interest should therefore include not only regulation of downstream harms, but also governance of design conditions, infrastructures, and socio-technical architectures. This connects to broader traditions of responsible innovation and value-sensitive design, which emphasise embedding democratic values and public accountability into technological systems from the outset (van den Hoven et al., 2015), while bringing explicitly into scope the life cycle, socio-technical, and infrastructural elements of a technology.

Public-interest neuro-AI infrastructures are not only objects of regulation, but also potential public goods.[2] If developed under appropriate conditions of

[2] The public goods argument requires qualification. Unlike paradigmatic public goods, in the case of neurodata, non-rivalry and non-excludability are not intrinsic properties but outcomes of institutional design: proprietary infrastructures, licensing regimes and data trusts can render neurodata excludable regardless of its technical shareability. Moreover, neurodata are extracted from and remain tied to identifiable bodies, carrying sensitive inferential potential (e.g. regarding cognitive states, affective dispositions, or neurological decline). This distinguishes neurodata from other, less identifiable forms of public data (such as weather or open government statistics). We use 'public good' here in a qualified sense, close to a commons requiring active stewardship.

transparency, scientific validity, and democratic control and accountability, they may contribute to advances in neuroscience, healthcare, education, and scientific discovery. Treating such infrastructures as public goods means recognising that decisions about their design, governance, access, and acceptable uses should not be determined solely by commercial incentives or technological feasibility. Instead, they should be subject to democratic deliberation and public accountability, with governance arrangements that ensure responsiveness to societal priorities, meaningful public participation, and independent oversight.

Properly governed, public neuro-AI infrastructures can create shared scientific resources, strengthen European research capacity, improve reproducibility and transparency, and enable the development of neuro-AI systems that advance collective capabilities and reflect democratic values and the public interest rather than exclusively commercial or geopolitical imperatives.

Governance should thus remain proportionate and ensure that it does not unnecessarily obstruct legitimate research, assistive neurotechnologies, accessibility-oriented innovation, or responsible public-interest neuroscience. This reflects broader European traditions of proportionate governance and rights-balancing under conditions of scientific and societal uncertainty (März et al., 2026). The aim is not to halt neurotechnology development, but to ensure that emerging neuro-AI infrastructures develop under conditions consistent with scientific validity, fundamental rights and the public interest, and democratic accountability. As such, this is compatible with the European commitment to responsible, societally-beneficial innovation.

The ethical and governance concerns discussed in this section thus arise at three interconnected levels: the individual, the collective, and the infrastructural. Neuro-AI systems may affect individual rights, shape group-level classifications and social outcomes, and contribute to the development of infrastructures whose effects extend across sectors and over time.

5. Regulatory interfaces and governance gaps

Existing regulatory frameworks have been developed to govern identifiable products, services, or uses. As we have explored, neuro-AI infrastructures challenge this approach because the same data, models, and capabilities may be reused, recombined, and redeployed across multiple contexts over time.

There is growing awareness across several jurisdictions of the need for regulatory action to keep pace with neurotechnological developments. US lawmakers' concerns over apparent lacunae in legal provisions over neurodata have led four US states to introduce neurotech regulations (Spivack and Victory, 2025; US Senate Committee on Commerce, Science & Transportation, 2025). Canada has decided that neurodata will be classified alongside other categories of sensitive information that require a higher degree of protection (OPC, 2026).

5.1. Data protection and neurodata-derived inferences

Neuro-AI infrastructures cut across several existing EU regulatory regimes. Under the General Data Protection Regulation (GDPR), key issues include the classification of neurodata and neurodata-derived inferences, the identification of a lawful basis for processing, requirements relating to special-category data and secondary use, consent limitations, rights of data subjects, consideration of collective effects, and requirements relating to third-country transfers (CCNE and CCNEN, 2026; CNIL, 2025; EDPS, 2024; Rainey et al., 2020). Without prejudice to the provisions of the GDPR, new legislation emerging as part of the EU's digital strategy, including the Data Act, the Data Governance Act and the Regulation on the European Health Data Space, is laying down (often sector-)specific rules.

The issue is not only whether neurodata constitute personal data (as defined by the GDPR), but also how derived inferences and model-training reuse are (and should be) governed. This is especially important where apparently low-grade or non-medical neurodata are combined with behavioural, contextual, or biometric data to generate sensitive inferences about mental, cognitive, affective, behavioural, or health-related states.

A further issue is that individual rights may be difficult to exercise effectively where neurodata are processed through complex infrastructures. Transparency, access, erasure, and restriction of processing become harder to operationalise when data have been aggregated, used for model training, combined with other datasets, or transformed into derived inferences.

Under the GDPR, data are classified according to the nature of the data or semantic content. Following a position taken by the French data protection authority CNIL (Commission National de l'Informatique et des Libertés), an important distinction can be made between health data "by nature" and health data "by destination". A "wellbeing' app recording heart rhythms was determined to be collecting health data rather than non-health or "wellness" data since it could be used to predict arrhythmia (CNIL, 2018). The same would apply to a smart watch that records heart rate variability. This determination makes the data collected by the app more sensitive than they would have otherwise been, and therefore subject to extra consideration and protection. As a minimum, this affects storage requirements for the data, which (in the context of French law) should be kept locally on the collection device.

The purposes for which direct-to-consumer applications of neuro-AI might be put include optimising meditation practices, self-monitoring, and mood management, among others. These look like parallels to the putative "wellbeing" category that CNIL has held corresponds to health data. It may not always be clear that all neurodata can easily be seen as health data at the point of collection, given the wide-ranging purposes for which brain recording and processing might occur, including activities such as gaming, flying drones, and typing on a PC. The following characteristics of these applications complicate matters even more. It is most likely that, given size, weight, and heat requirements, neuro-AI devices would require cloud storage and data processing capacities to enable devices to remain small, light, and cool. Data processing and transfer should therefore be considered as built-in features of such technologies. Additionally, the training of AI requires large amounts of data, so it must be assumed that part of the process of remote storage of data in these devices will include data aggregation as a means of increasing the pool available for training machine learning algorithms. From a device manufacturer's perspective, neurodata in such cases constitute an asset to boost AI usability.

Still, given the infrastructural and public interest perspective foregrounded in this Statement and following the analogy suggested by the CNIL, we argue that neurodata collected by consumer neurotech applications aiming at leisure, wellness, sleep, lifestyle, or cognitive enhancement purposes should count as health data and ought to be governed and regulated as such.

As long as there is a question mark over the purpose and future use of the data collected – and in light of the fact that such data will be infrastructural to this kind of technology, remote storage and processing will be a norm, and aggregation will be a part of the business model of neuro-AI companies – the GDPR dimensions of neuro-AI technologies must be carefully considered (Rainey et al., 2020). If neurodata reveal information about a person's health, mental states, or neurological condition, they are clearly "data concerning health" under Article 4(15) and Article 9(1) of the GDPR. In such cases, they are classified as sensitive, and processing them is prohibited unless an Article 9(2) exemption and a legal basis according to Article 6(1) GDPR applies.

It is arguable that neurodata should be explicitly included in the category of "special category" data in Article 9 of the GDPR, thus prohibiting their processing unless additional protections and exemptions are met. Existing categories may cover many instances of neurodata processing, but uncertainty remains regarding consumer applications, downstream inference, and foundation-model training. Explicit reference to neurodata in the definition would remove such uncertainty. The implications of such a revision of Article 9(1) of the GDPR, which opens up the possibility of applying further relevant provisions, should be taken into account to provide guidance to those applying the law and to ensure consistency in its application. For example, where data processors rely on the exception for scientific research purposes set out in Article 9(2)(j) of the GDPR, they must carry out a necessity and proportionality assessment, for which they will in all likelihood require yardsticks.

Another concern is that the applications on which we have focused here could be said to aim for something other than the unique identification of natural persons and related processing purposes. They require data for the training of large BFMs/HFMs that can later help to classify and profile individuals (not as individuals but as members of a distinct category of individuals) and predict and make determinations concerning further downstream uses based on a pre-trained model. By explicitly including neurodata in a "special category" of data according to Article 9 GDPR, it could also be suggested that further conditions for its processing should be laid down at European legislative level, or that the legislator should guide the relevant balancing of competing rights and interests, while maintaining fairness in single-case law application.

Due to the special characteristics of neurodata and the particularities of their processing, the control options available to data subjects are also rather limited. To exercise the individual rights set out in Chapter III of the GDPR, not only is comprehensive information about the processing necessary, but it must also be ensured that the data subject can respond to the processing directly and consciously. Given its direct connection to the cognitive system of the data subject, this is not possible in many cases of neurodata processing without additional information.

Particular attention should be paid to compliance with information obligation requirements, and specific solutions – such as automated access to or retrieval of information – should be considered.

5.2. The AI Act and foundation-model governance

Various sections of the AI Act are potentially relevant to the use of neurodata-based AI applications, starting with the prohibitions on unacceptable-risk AI, such as harmful subliminal techniques, manipulation, deception, exploitation of vulnerabilities, social scoring with detrimental effects, emotion recognition in workplaces and educational settings or certain biometric categorisation systems (see below). Other sections of relevance include those relating to high-risk systems, general-purpose AI, foundation-model governance and fundamental-rights impact assessment, among many. However, while some neuro-AI applications and neurotech devices might currently be placed with some confidence in the “unacceptable risk AI systems” category (or in the “high risk” category if purely medical applications are concerned), it will often be unclear how to classify others under the law as it stands. Applications such as consumer brain recording and neurofeedback devices, for example, might be seen as low risk in only recording quite crude neural signals, without much genuine capacity for indicating or affecting users’ cognitive, emotional, or other mental states, or uniquely identifying them on that basis.

As we have argued, neuro-AI systems challenge existing regulatory categories because their risks do not arise solely from the intended purpose of individual devices or applications. Even low-grade neural signal recording of the type envisaged in the “low-risk” example above may contribute valuable training data to reusable BFMs/HFMs, thereby supporting the development of increasingly sophisticated neuro-AI capabilities. Over time, neural data collected across multiple contexts and devices may be aggregated and analysed for purposes that extend well beyond those originally envisaged at the point of collection, including applications that have yet to be conceived.

The principal governance challenge therefore lies in recognising that reusable infrastructures generate capabilities whose downstream applications may evolve over time. Existing regulatory approaches primarily assess identifiable systems placed on the market for specified purposes. Neuro-AI infrastructures instead accumulate capabilities that may subsequently be recombined across multiple sectors, requiring governance that accompanies capabilities throughout their lifecycle rather than only their initial deployment.

This poses an important limitation within the current regulatory architecture. While the AI Act represents a significant step forward by recognising the importance of regulating general-purpose AI models, its risk classification remains largely focused on specific applications and intended uses. Neuro-AI infrastructures, by contrast, are characterised by downstream modularity, reconfiguration, and cross-contextual deployment. A model initially developed for research or wellness purposes may subsequently underpin applications in healthcare, employment, education, insurance, security, or defence. **Governance should therefore extend beyond the assessment of individual (types of) devices or applications to encompass the reusable model and infrastructure layer, where cumulative capabilities and**

societal risks increasingly emerge. Such an approach would not imply that all brain-recording devices should be classified as presenting unacceptable risk. Rather, it recognises that seemingly low-risk systems may contribute to infrastructures whose downstream uses require ongoing oversight, accountability, and democratic governance.

This infrastructural perspective also has implications for interpreting the AI Act. Neuro-AI infrastructures may, through large-scale neural data aggregation, multimodal data fusion, and increasingly-capable foundation models, support downstream applications that fall within existing prohibited or high-risk categories, including manipulation, exploitation of vulnerabilities, emotion recognition, biometric categorisation, or other practices that interfere with fundamental rights. The principal governance challenge, however, is that these capabilities emerge cumulatively through reusable models, large-scale neurodata aggregation, continual adaptation, and cross-context deployment.

This exposes an important limitation of predominantly application-based regulation. Although the AI Act recognises the importance of regulating identifiable AI systems and general-purpose AI models, its core architecture remains centred on identifiable AI systems, intended uses, and risk categories. Like other advanced AI infrastructures, neuro-AI infrastructures may generate capabilities whose future applications may not be foreseeable at the time of development or market placement. However, this challenge is amplified in neuro-AI by the combination of neural data aggregation, multimodal integration, and the potential sensitivity of applications affecting cognition, behaviour and identity. Governance should therefore complement application-specific regulation with infrastructure-level oversight, including requirements concerning provenance, documentation, monitoring, transparency, accountability, and lifecycle governance of reusable neuro-AI infrastructures.

5.3. The Medical Device Regulation and intended-use problems

The Medical Device Regulation (MDR) raises additional challenges where consumer-facing neurotechnologies generate quasi-clinical effects or health-related inferences while remaining outside the MDR's scope due to the absence of a declared medical purpose. The issue is arguably not the manufacturer's declared medical intention but the practical function and social effects of systems that shape how users understand and act upon information about their health, cognition, emotional state, or mental wellbeing.

Products marketed as tools for sleep optimisation, wellness, mood management, stress reduction, productivity, or cognitive enhancement may still produce inferences that users interpret as medically or psychologically meaningful. In practice, such systems can influence users' perceptions of their neurological or mental condition, encourage behavioural changes, or affect decisions about seeking treatment or self-management, even where no explicit medical claim is made. Intended-use classification may therefore be insufficient in cases where products are technically framed as non-medical but function "by destination" or by practical effect in ways that approximate clinical assessment or intervention.

This challenge resembles broader debates in data protection law, including discussions associated with the French data protection authority's notion of functioning “by destination” where the practical use and societal effects of a technology may become more significant than its formally declared purpose (see section 5.1 above) (CNIL, 2018). Similar tensions arise under both the GDPR and the MDR whose formal classifications and declared purposes may fail to capture systems that generate highly sensitive inferences or quasi-clinical consequences in practice including health-related recommendations, cognitive enhancement, diagnostic-like outputs, or therapeutic effects (Ienca et al., 2022; Rainey et al., 2020; Svempe, 2025).

5.4. Consumer protection, product safety, and digital fairness

Similar concerns arise under product safety, consumer protection, and digital fairness frameworks, particularly where neuro-AI products involve overclaiming, weak scientific validity, unfair terms, dependency creation, dark patterns (i.e. design choices that steer users into making decisions they would not otherwise make, typically in ways that benefit the service provider rather than the user), or opaque or misleading claims about mental or cognitive benefits (CCNE and CCNEN, 2026; Timms, 2025).

Consumer protection is especially important where users are not well positioned to assess the scientific validity of claims, for example, about sleep, attention, stress, mood, cognition, or mental performance. It is also important where the user-facing service obscures downstream data uses, model training, or third-party access.

For example, a wearable marketed as a non-medical safety or productivity tool may classify attention, fatigue, or stress. Even if framed as workplace support, such classifications may potentially affect evaluation, employment conditions, insurance, liability, or disciplinary decisions in professional contexts (Wascher et al., 2014).

5.5. Cross-border processing, “dual use”, and regulatory arbitrage

Neuro-AI infrastructures raise additional concerns regarding cross-border processing, third-country transfers, security and military use, and regulatory arbitrage. Systems may be trained, hosted, updated, or repurposed across jurisdictions, including jurisdictions with varying levels of data protection, consumer protection, labour protection, or democratic oversight.

This is particularly relevant where neurodata collected in consumer or research contexts are later used to train models deployed in commercial, workplace, security, or military settings. In the near future, governance therefore needs to address not only point-of-collection compliance, but also downstream reuse, transfer, deployment, and institutional accountability.

Regulatory arbitrage may also occur where actors strategically frame systems as wellness, entertainment, productivity, or research tools to avoid stricter medical, AI, or consumer-protection obligations. A targeted gap analysis should therefore examine not only formal classifications, but also foreseeable effects, downstream uses, and infrastructural dependencies.

5.6. Socially responsible innovations in neurotechnology

Neurotechnology and AI – severally and combined – hold promise for improving human health and well-being. Some of the most relevant examples would be the remarkable results achieved by cortical neurotechnology to help recover sensory-motor function as in the treatment of severe Parkinson's disease, and through neuromodulation in neuropsychiatry to help manage debilitating conditions. A deeper understanding of the working of the human brain can undoubtedly lead to fruitful insights that improve human health and wellbeing.

But the significant ethical challenges associated with the development of these technologies require careful consideration and robust regulatory frameworks that also address the infrastructural dimension. These are essential requirements for socially responsible innovation – that is, innovation that contributes to the common good and is not held hostage to the economic and financial interests of a few. Protecting human rights and human dignity in the face of these advancements is paramount to ensuring that neuro-AI benefits both individuals and societies.

6. Recommendations

The European Group on Ethics in Science and New Technologies recommends that European values – enshrined in the constitutive and binding EU instruments on fundamental and human rights – should drive responsible innovations in neurotechnology.

We aim to contribute to preparing the EU for an appropriate anticipatory governance posture and adequate regulation, thereby promoting such responsible innovations. Neuroscience and neurotechnology powered by developments in Brain Computer Interface (BCIs: direct interface between the human brain and an external device), big neurodata, Brain Foundation Models and Human (centric) Foundation Models (BFMs/HFMs: large-scale AI models trained on neural, neuroimaging, behavioural or contextual data, and adaptable across downstream tasks), digital twins, deep learning, and generative and agentic AI pushes us to the boundaries in our thinking about privacy, freedom, autonomy, human dignity, and what it means to be human.

It also requires the EU to develop governance approaches for an emerging neurostack, a coordinated ensemble of devices, data pipelines, BFMs/HFMs, cloud services, interfaces, business practices, standards, audit and oversight mechanisms and institutional arrangements through which neurodata are collected, processed, reused, and transformed into practical applications. It should be designed and governed under EU law and values to ensure scientific validity, fundamental-rights compliance, security and privacy by design, public interest capacity, and digital sovereignty across the lifecycle of neuro-AI systems. The governance of emerging neuro infrastructure would emphasise European democratic control and accountability in shaping the common cognitive and epistemic infrastructure on which future neuro-AI applications depend.

In this context, the European Commission should evaluate the capacity of existing regulatory frameworks to govern emerging AI–neurotechnology integrations, while identifying regulatory gaps arising from novel combinations of neurodata, foundation models, and adaptive computational infrastructures.

The capacity of market actors to adapt rapidly to evolving regulatory environments highlights the need for a more coherent governance approach. Regulation concerning medical devices, consumer protection, data protection, and prohibited or high-risk AI systems should not operate in isolation, but rather as complementary layers of oversight capable of addressing the cross-sectoral nature of neuro-AI. Such coordination is necessary to ensure that neuro-AI technologies develop in ways that serve fundamental rights, individual autonomy, and the European and global public interest.

The five recommendations set out below operate at different levels of governance. For each recommendation, we have identified the relevant governance objective. Some objectives focus on the elucidation and application of existing legal frameworks, while others address broader questions of institutional capacity, infrastructure governance, and strategic oversight. Taken together, they are intended to provide a coherent governance approach spanning the full lifecycle of neuro-AI systems, from data collection and model design and development to deployment, societal impact, and public accountability.

The recommendations are organised around five complementary governance objectives:

- Recommendation 1 – Protection of neurodata: Define and reinforce protections for neurodata and neurodata-derived inferences under EU data protection law.
- Recommendation 2 – Responsible development of neuro-AI technologies: Ensure possible development of Brain Foundation Models and related neuro-AI infrastructures.
- Recommendation 3 – Rights protection: Ensure that individuals are protected against disproportionate control where neurodata and neurodata-derived data inferences are used in consequential contexts.
- Recommendation 4 – Governance capacity: Build European public-interest governance capacity for neuro-AI infrastructures.
- Recommendation 5 – Fitness check: Conduct a targeted regulatory assessment of the EU framework governing neuro-AI systems.

Recommendation 1: Neurodata and neurodata-derived inferences should explicitly be listed as a special category of sensitive data under EU data protection law

The European Commission, Member States, and relevant Union bodies should govern neurodata and neurodata-derived inferences as a distinct category of sensitive data, considering both individual and wider societal implications, and develop coordinated regulatory guidance and, where appropriate, legislative clarification regarding the legal qualification, processing, and overall governance of neurodata and neurodata-derived inferences. This should include clarification of the circumstances under which such data should be subject to heightened safeguards equivalent or comparable to those applicable to the special categories of personal data across EU data protection law and in the legal frameworks governing consumer protection, digital services, product safety, medical devices, and the use of AI. Such EU guidance should also clarify that neurodata-derived inferences may warrant protection equivalent to the underlying neural data where they reveal information about cognitive states, mental processes, health, behaviour, preferences, or vulnerabilities.

The EU should develop a dedicated **Neuro-AI impact assessment methodology** for systems processing neurodata or generating neurodata-derived inferences, addressing risks to privacy, autonomy, cognitive integrity, discrimination, and collective harms.

Such governance should not be limited to individual data points or requirements for individual consent. While valid consent may provide a lawful basis for certain processing activities, it is not sufficient to address the broader risks arising from the collection, analysis, combination, and downstream use of neurodata, owing to its impact on individuals beyond those providing consent, groups, and society as a whole.

EU governance should address the full lifecycle of neurodata, including collection, annotation, aggregation, model training, inference generation, secondary use, and cross-context deployment.

More specifically, governance should address how neurodata and neurodata-derived inferences acquire meaning, predictive value, and economic utility through aggregation, comparison, modelling, and combination with other datasets, and how their reuse may create risks also for individuals, groups and society at large, beyond the original data subject. This includes secondary use, cross-context transfer, training-data reuse, and collective profiling consistent with GDPR provisions on profiling and behavioural prediction (Article 4(4) and Article 22), the principles of data minimisation (Article 5(1)(c)) and purpose limitation (Article 5(1)(b), and emerging EU case law and regulatory guidance on inferential data and high-risk AI systems under the AI Act.

The European Commission, together with the European Data Protection Board (EDPB), the European Data Protection Supervisor (EDPS), and relevant scientific bodies, should develop an EU-level taxonomy of neurodata and neurodata-derived inferences, distinguishing between raw neural signals, processed neural features, inferred cognitive or emotional states, and models trained on neurodata.

Recommendation 2: Ensure responsible development of BFMs/HFMs

The European Commission, Member States, and relevant Union bodies should develop, or require the development of, technical standards, regulatory guidance, and harmonised governance requirements for BFMs and related neurodata infrastructures. This work should be coordinated, where appropriate, with European and international standardisation bodies and conformity-assessment mechanisms to promote interoperability, reproducibility, scientific validity, auditability, and consistent implementation across the Union.

As a minimum, these governance requirements should address the following:

- the need to ensure that neurodata used in BFMs/HFMs have a documented and auditable provenance chain and are processed in compliance with Regulation (EU) 2016/679 (GDPR) including the identification of an appropriate lawful basis for the collection and use of training data;
- consent conditions, where applicable, that are explicit, informed, and granular (the prerogative to grant or withhold consent for specific, individual uses of personal data, rather than providing blanket approval for all potential uses), accompanied by recognition of the limits of individual consent as a governance mechanism for large-scale model development and updating and downstream reuse;
- cross-context reuse, access and release conditions, including the development of model documentation;
- downstream deployment, including the use of a dual-use risks inventory that covers systematic assessment and documentation of dual-use risks, misuse scenarios, and potential impacts on fundamental rights, public security, and democratic processes;
- third-country transfers that comply with GDPR Chapter V requirements;
- adequate safeguards for adequacy decision and proportionality;
- fair benefit sharing mechanisms; and
- safeguards against regulatory arbitrage.

Particular attention should be paid to models trained on neurodata originating from consumer neurotechnology applications that may be subject to weaker governance standards, including commercial direct-to-consumer neurotechnology products and services used for sleep, wellness, cognitive-enhancement, mood-management, gaming, or workplace applications.

Recommendation 3: Ensure that individuals are protected against disproportionate control where neurodata or neurodata-derived inferences are used in the exercise of decision-making powers on the part of either public or private actors

The EU should establish clear prohibitions and safeguards to prevent the use of neuro-AI systems to influence, manipulate, coerce, exploit vulnerabilities, or classify, rank or profile persons or groups in ways incompatible with human dignity, autonomy, privacy, non-discrimination, freedom of thought, bodily integrity, or mental integrity. All prohibitions and safeguards shall ensure that individuals are protected against disproportionate epistemic or institutional control resulting from the collection, processing, analysis, inference, or use of neurodata in the exercise of decision-making powers by public or private actors. Particular scrutiny should apply to the deployment of neuro-AI systems in areas such as employment, education, healthcare, insurance, law enforcement, migration, border control, criminal justice, social protection, and public services.

Safeguards should include requirements for technical robustness, proportionality, scientific validity, contextual justification, meaningful human oversight and review, contestability, independent audits, transparency about uncertainty and limits of inference, and restrictions on high-impact uses.

Particular attention should also be paid to systems that influence behaviour through optimisation processes, including the optimisation of engagement, learning, productivity, compliance, or health outcomes.

Recommendation 4: Build European public-interest governance capacity for neuro-AI infrastructures

The European Commission should, in coordination with Member States and relevant Union agencies, strengthen the capacity of the EU to monitor, assess, and shape neuro-AI infrastructures in the public interest in alignment with Union law, including the Charter of Fundamental Rights of the European Union, the General Data Protection Regulation (GDPR), the EU AI Act and related digital and health regulatory frameworks.

This should include monitoring, horizon-scanning, and strategic foresight functions; independent standards, certification schemes, audit and conformity-assessment mechanisms; fundamental-rights impact assessment and post-market monitoring; public-interest research capacity; public procurement criteria; regulatory sandbox initiatives under robust rights-based and safety conditions; and structured public and expert deliberation. Existing or new observatory functions should avoid institutional duplication and be clearly embedded in EU governance structures, with defined mandates, interoperability with national competent authorities, and clear lines of accountability at Union level.

The aim should be to align neuro-AI development with the Union acquis concerning scientific validity, public health, democratic accountability, fundamental rights, data protection, privacy-by-design and security-by-design principles, and responsible innovation, while reducing risks of market concentration, data capture, and reducing technological dependence on non-EU platform infrastructures.

In governing an emerging neurostack, particular attention should be given to systemic risks, including those related to high-impact foundation models, neurotechnological interfaces, and cross-border data governance asymmetries, in line with the EU's digital sovereignty, technological resilience, and strategic autonomy objectives including the preservation of European values in the development and deployment of neuro-AI systems.

Recommendation 5: Conduct a fitness check (targeted regulatory assessment) of the EU framework governing neuro-AI systems

The European Commission should initiate in cooperation with the AI Office, the European Data Protection Board (EDPB), or the European Data Protection Supervisor (EDPS), and other relevant Union bodies, a targeted assessment of the adequacy of EU law (including the GDPR, AI Act, Medical Device Regulation, product safety law, consumer protection law) to govern neuro-AI systems, including direct-to-consumer neurotechnology and Brain Foundation Models/Human (centric) Foundation Models, irrespective of whether such systems are currently classified as medical devices, consumer products, or high-risk AI systems under existing frameworks.

This review should assess the coherence, effectiveness, and suitability of the current regulatory framework and identify potential gaps, inconsistencies, or enforcement challenges relating to intended-use classification, non-medical but high-impact applications, neurodata-derived inferences, collective, societal, and systemic harms, dual use risks, impact on vulnerable groups, third-country processing, and benefit sharing.

It should further assess the adequacy of existing safeguards under Union law, including data protection, consumer protection, product safety, and fundamental rights frameworks, for systems that combine neurodata, behavioural data, foundation-model architectures, and downstream inferential uses, taking into account the risks such systems may pose to individuals, groups, and society more broadly.

The fitness check should also assess the capacity of the current EU framework to remain effective and adaptable in light of the rapid evolution of neuro-AI technologies, including emerging architectures, changing deployment models, and novel forms of data-driven inference.

ANNEX 1

Key Terms

Neurodata: Data relating to the structure, activity, or function of the nervous system, including data collected through neurotechnology or neurophysiological sensing.

Neurodata-derived inference: A prediction, classification, assessment, or other conclusion generated from neurodata alone or in combination with behavioural, biometric, contextual, or health-related data. Such inferences may be generated not only from direct neural measurements, but also from combinations of neurophysiological, behavioural, biometric, and contextual data.

Brain Foundation Model: BFMs are AI models pretrained on large, heterogeneous, and often un-labelled neuroimaging datasets spanning electrophysiological, functional, and structural modalities, including EEG, MEG (magnetoencephalography), fMRI (functional magnetic resonance imaging), MRI, and PET (positron emission tomography).

Human (centric) Foundation Model: A human foundation model (also known as a human-centric foundation model) is an advanced AI system trained on massive datasets of human data – such as body geometry, facial expressions, speech, and movement. These versatile systems serve as a baseline for creating specialised digital avatars, robotic systems, and generative content.

Neuro-AI infrastructure: The wider system of devices, data pipelines, models, cloud services, interfaces, business practices, and institutional arrangements through which neurodata are collected, processed, reused, and operationalised.

Neurostack: A coordinated ensemble of devices, data pipelines, BFMs/HFMs, cloud services, interfaces, business practices, standards, audit and oversight mechanisms, and institutional arrangements through which neurodata are collected, processed, reused, and operationalised.

Consequential Context: A setting in which neural or behavioural inferences may affect a person's rights, opportunities, access to services, employment, education, insurance, healthcare, legal standing, security treatment, or social status.

ANNEX 2

References

Alegre, S., Freedom to think: The long struggle to liberate our minds, Atlantic Books, 2021.

Anzolin, E. and Lo Nostro, G., 'Ray-Ban Meta glasses take off but face privacy and competition test',Reuters website, 9 December 2025, accessed 30 June 2026, https://www.reuters.com/sustainability/boards-policy-regulation/ray-ban-meta-glasses-take-off-face-privacy-competition-test-2025-12-09/.

BCC Research, Neurotech devices: Global market outlook, BCC Publishing, 2023.

Biospace, 'Neurotechnology market size to surpass USD 52.86 billion by 2034 driven by breakthroughs in brain-machine interfaces', Biospace website, 2026, accessed 24 June 2026, https://www.biospace.com/press-releases/neurotechnology-market-size-to-surpass-usd-52-86-billion-by-2034-driven-by-breakthroughs-in-brain-machine-interfaces.

Boonstra, J.T., 'Ethical imperatives in the commercialization of brain-computer interfaces', IBRO Neuroscience Reports, Vol. 19, 2025, pp. 718-724, https://www.ibroneuroreports.org/article/S2667-2421(25)00160-5/pdf.

Bowker, J. and Star, S.L., Sorting things out: Classification and its consequences, MIT Press, 1999.

Bublitz, J. C., 'Freedom of thought in the age of neuroscience: A plea and a proposal for the renaissance of a forgotten fundamental right', Archiv für Rechts- und Sozialphilosophie / Archives for Philosophy of Law and Social Philosophy, Vol. 100(1), 2014, pp. 1–25.

CCNE and CCNEN, Avis 150 du CCNE et Avis 10 du CCNEN sur les neurotechnologies numériques, Comité consultatif national d'éthique / Comité consultatif national d'éthique du numérique, 2026, https://www.ccne-ethique.fr/sites/default/files/2026-07/AVIS-COMMUN_CCNE_150_CCNEN_10.pdf.

Centre for Future Generations, Neurotech between wellness and medicine: Towards an EU neurotechnology strategy, Centre for Future Generations, 2026.

CNIL, Qu'est-ce ce qu'une donnée de santé?, CNIL, 2018, https://www.cnil.fr/fr/quest-ce-ce-quune-donnee-de-sante.

CNIL, Neurotechnologies, from the realm of possibilities to regulatory projects;

Neurodata, personal data like no other, CNIL, 2025, https://linc.cnil.fr/en/neurotechnologies-realm-possibilities-regulatory-projects.

Constine, J., Facebook buys CTRL-Labs for between $500M and $1B, TechCrunch, 2019.

DARPA, ‘N3: Next-generation nonsurgical neurotechnology’, DARPA website, undated, accessed 5 June 2026, https://www.darpa.mil/research/programs/next-generation-nonsurgical-neurotechnology.

Dey, M., Dastin, H. and Thomas, C., ‘Anthropic disables top-tier AI models after US order limiting foreign access’, Reuters website, 13 June 2026, accessed 25 June 2026, https://www.reuters.com/technology/us-blocks-foreign-access-anthropics-most-advanced-ai-models-axios-reports-2026-06-13/.

Dibble, A., Dalby, C., Sevegnani, M., Fracasso A., Lyall, D.M. et al., ‘NeuroFM: Toward precision neuroimaging with foundation models for individualized brain health estimation’, medRxiv: 2026.03.27.26349489, 2026, https://www.medrxiv.org/content/10.64898/2026.03.27.26349489v2.full.pdf.

European Data Protection Supervisor (EDPS), TechDispatch #1/2024 – Neurodata, 2024, https://www.edps.europa.eu/data-protection/our-work/publications/techdispatch/2024-06-03-techdispatch-12024-neurodata_en.

European Group on Ethics in Science and New Technologies (EGE), Statement on artificial intelligence, robotics and 'autonomous' systems, European Commission, 2018, https://ethicsadvice.eu/advice/artificial-intelligence-robotics-and-autonomous-systems/.

European Group on Ethics in Science and New Technologies (EGE), Democracy in the Digital Age, European Commission, 2023, https://ethicsadvice.eu/advice/democracy-in-the-digital-age/.

European Parliament, ‘EU AI act: First regulation on artificial intelligence’, European Parliament website, 2023, accessed 30 June 2026, https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence.

Finlay, L. and Hooten, P., ‘When war reaches the human mind, Australia must set the rules’, The Strategist website, 24 April 2026, accessed 5 June 2026, https://www.aspistrategist.org.au/when-war-reaches-the-human-mind-australia-must-set-the-rules/.

Geneva Academy, ‘Anticipating human rights risks in neurotechnology: new guidance for states and businesses’, Geneva Academy website, 29 January 2025, accessed June 30, 2025, https://geneva-academy.ch/anticipating-human-rights-risks-in-neurotechnology-new-guidance-for-states-and-businesses/.

Genser, J., Damianos, S. and Yuste, R., Safeguarding brain data: Assessing the privacy practices of consumer neurotechnology companies, Neurorights Foundation, 2024, https://perseus-strategies.com/wp-content/uploads/FINAL_Consumer_Neurotechnology_Report_Neurorights_Foundation_April-1.pdf.

Gielas, A., 'Warfare at the speed of thought: Can brain-computer interfaces comply with IHL?', ICRC website, 21 August 2025, accessed 5 June 2026, https://blogs.icrc.org/law-and-policy/2025/08/21/warfare-at-the-speed-of-thought-can-brain-computer-interfaces-comply-with-ihl/.

Github, 'Awesome neuro-foundation models', Github website, undated, accessed 3 June 2026, https://github.com/mazabou/awesome-neurofm.

Global Privacy Assembly, Resolution on neurotechnology, human rights, data protection and privacy, 2024. https://globalprivacyassembly.com/wp-content/uploads/2024/11/Resolution-on-Neurotechnologies.pdf.

Greenbaum, D., 'Ethical, legal and social concerns relating to exoskeletons', ACM SIGCAS Computers and Society, Vol. 45(3), 2016, pp. 234-239, https://dl.acm.org/doi/epdf/10.1145/2874239.2874272.

Guenther, S., Kosmyna, N. and Maes, P, 'Image classification and reconstruction from low-density EEG', Scientific Reports, Vol. 14, 2024, p. 16436, https://www.nature.com/articles/s41598-024-66228-1.

Haar Horowitz, A., Stickgold, R. and Zadra, A., 'Inside your dreamscape', Aeon website, 2021, accessed 30 June 2026, https://aeon.co/essays/dreams-are-a-precious-resource-dont-let-advertisers-hack-them.

Hanley, M., Yeh, J.-T., Rodriguez, R., Pilkington, J. and Farahany, N, 'Training data governance for brain foundation models', arXiv:2602.02511, 2026, https://arxiv.org/abs/2602.02511.

Helbing, D. and Ienca, M., 'Why converging technologies need converging international regulation', Ethics and Information Technology, Vol. 26(1), 2024, p. 15, https://link.springer.com/article/10.1007/s10676-024-09756-8.

Horikawa, T., Tamaki, M., Miyawaki, Y. and Kamitani, Y., 'Neural decoding of visual imagery during sleep', Science, Vol. 340(6132), 2013, pp. 639–642.

Ienca, M., Fins J.J., Jox R.J., Jotterand, F., Voeneky, S. et al., 'Towards a governance framework for brain data', Neuroethics, Vol. 15(2), 2022, p. 20, https://link.springer.com/article/10.1007/s12152-022-09498-8.

Kane, W., 'Neurocapitalism: Dream enclosure, neural rent, and the final governance of interior life', 2025. https://philarchive.org/rec/KANNDE.

Kang, Q., Liu, F. and Cai, C., 'Effective and efficient self-supervised masked model based on mixed feature training', Frontiers in Neurorobotics, Vol.19, 2025, p. 1705970,https://www.frontiersin.org/journals/neurorobotics/articles/10.3389/fnbot.2025.1705970/full.

Kellmeyer, P., Biller-Andorno, N. and Meynen, G., 'Ethical tensions of virtual reality treatment in vulnerable patients', Nature Medicine, Vol. 25(8), 2019, pp. 1185–1188.

Kwon, J. and Shin, Y., 'Foundation models for neural signal decoding: EEG-centered perspectives toward unified representations', European Journal of Neuroscience, Vol. 63(1), 2026, e70376.

Lepapa, N., Abdigadir, A. and Lindblom, J., 'She came out of the bathroom naked, employee says', Svenska Dagbladet website, 27 February 2026, accessed 30 June 2026, https://www.svd.se/a/K8nrV4/metas-ai-smart-glasses-and-data-privacy-concerns-workers-say-we-see-everything.

Mantellassi, F. and Madziwa, E., Neurotechnology in the military domain: A primer, UNIDIR, 2025, https://unidir.org/wp-content/uploads/2025/11/UNIDIR_Neurotechnology_Military-Domain_A-Primer.pdf.

Marlan, D., 'The nightmare of dream advertising', William & Mary Law Review, Vol. 65(2), 2023, pp. 259-326, https://scholarship.law.wm.edu/cgi/viewcontent.cgi?article=4005&context=wmlr.

März, J., Biller-Andorno, N., Mouton Dorey, C. and Dagron, S. (Eds.), The principle of proportionality in public health law, ethics, and policy, Oxford University Press, 2026, https://academic.oup.com/book/61863.

Meta, 'Introducing tribe v2: A predictive foundation model trained to understand how the human brain processes complex stimuli', Meta website, 26 March 2026, accessed 24 June 2026, https://ai.meta.com/blog/tribe-v2-brain-predictive-foundation-model/.

Musk, E. (X:@elonmusk), 'You are your brain', 2 October 2025, accessed 26 June 2026, https://x.com/elonmusk/status/1974029419120414926.

National Institutes of Health (NIH), 'BRAIN initiative: Data knowledgebase ecosystem and neuroai integration', NIH Grants & Funding website, 24 February 2026, accessed 25 June 2026, https://grants.nih.gov/funding/find-a-fit-for-your-research/highlighted-topics/18.

OECD, Responsible innovation in neurotechnology enterprises, OECD Science, Technology and Industry Working Papers, 2019, https://www.oecd.org/content/dam/oecd/en/publications/reports/2019/10/responsible-innovation-in-neurotechnology-enterprises_2d346c46/9685e4fd-en.pdf.

OECD, Neurotechnology toolkit to support policymakers in implementing the OECD recommendation on responsible innovation in neurotechnology, 2025, https://www.oecd.org/content/dam/oecd/en/topics/policy-sub-issues/emerging-technologies/neurotech-toolkit.pdf.

Office of the Privacy Commissioner of Canada (OPC), 'Interpretation bulletin: Sensitive information', OPC website, 2026, accessed 30 June 2026, https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/pipeda-compliance-help/pipeda-interpretation-bulletins/interpretations_10_sensible/.

Olsen, S. G., Reinecke-Tellefsen, C. J. and Østergaard, S. D., 'Potentially harmful consequences of artificial intelligence (AI) chatbot use among patients with mental illness: early data from a large psychiatric service system', Acta Psychiatrica Scandinavica, Vol. 153(4), 2026, pp. 301–303, https://onlinelibrary.wiley.com/doi/epdf/10.1111/acps.70068.

OSCE Office for Democratic Institutions and Human Rights (ODIHR), Think again: Freedom of thought in the age of artificial intelligence, Organization for Security and Co-operation in Europe, 2023, https://odihr.osce.org/sites/default/files/documents/publications/2026/05/Freedom%20of%20Thought%20and%20AI%20-%20web%202025_10_3.pdf.

Pan, P., Tan, G., Wai, P. and Aung, A., 'Evaluation of consumer-grade EEG headsets for BCI drone control', Proceedings of the IRC Conference on Science, Engineering, and Technology, 2017, https://www.researchgate.net/publication/323692301_Evaluation_of_Consumer-Grade_EEG_Headsets_for_BCI_Drone_Control.

Pongiglione, F., 'Epistemic hubris', Social Epistemology, Vol. 39(1), 2025, pp. 91–105.

Prainsack, B., El-Sayed, S., Forgó, N., Szoszkiewicz, Ł. and Baumer, P., 'Data solidarity: A blueprint for governing health futures', The Lancet Digital Health, Vol. 4(11), 2022, e773-e774, https://www.thelancet.com/action/showPdf?pii=S2589-7500%2822%2900189-3.

Prainsack, B., States of solidarity: How to build a society. Oxford University Press, 2026.

Rainey, S., Philosophical perspectives on brain data, Palgrave Macmillan, 2023.

Rainey, S., McGillivray, K., Akintoye, S., Fothergill, T., Bublitz, C., and Stahl, B., 'Is the European Data Protection Regulation sufficient to deal with emerging data concerns relating to neurotechnology?', Journal of Law and the Biosciences, Vol. 7(1), 2020, lsaa051, https://academic.oup.com/jlb/article/7/1/lsaa051/5864051.

Reuben C., Elgaddal, N. and Black, L.I., Sleep medication use in adults aged 18 and over, National Center for Health Statistics, Centers for Disease Control and Prevention, 2023, https://www.cdc.gov/nchs/data/databriefs/db462.pdf.

Reuters, 'Meta expands AI access on Ray-Ban smart glasses in Europe', Reuters website, 23 April 2025, accessed 30 June 2026, https://www.reuters.com/business/meta-expands-ai-access-ray-ban-smart-glasses-europe-2025-04-23/.

Scientific Foresight Unit (STOA), The protection of mental privacy in the area of neuroscience: Societal, legal and ethical challenges, European Parliamentary Research Service, 2024, https://www.europarl.europa.eu/RegData/etudes/STUD/2024/757807/EPRS_STU(2024)757807_EN.pdf.

Spivack, J. and Victory, C., 'The "neural data" goldilocks problem: Defining "neural data" in U.S. state privacy laws', The Future of Privacy Forum website, 12 August 2025, accessed 30 June 2026, https://fpf.org/blog/the-neural-data-goldilocks-problem-defining-neural-data-in-u-s-state-privacy-laws/.

Stinchfield, B. T., 'The military and commercial development of brain–computer interfaces: International (in)security with brain-machine teaming', Defense & Security Analysis, Vol. 39(2), 2023, pp. 233–252.

Susser, D. and Cabrera, L. Y., 'Brain data in context: Are new rights the way to mental and brain privacy?' AJOB Neuroscience, Vol 15(2), 2024, pp. 122–133.

Svempe, L., 'The regulatory landscape of health apps in the European Union', JIPITEC – Journal of Intellectual Property, Information Technology and E-Commerce Law, Vol. 16(1), 2025, pp. 24-34, https://www.jipitec.eu/jipitec/article/view/419/422.

Tang, S., Wang, Y., Chen, L., Wang, Y. Peng, S. et al., 'Human centered foundation models: Perception, generation and agentic modeling', arxiv:2502.08556, 2025, https://arxiv.org/pdf/2502.08556.

Timms, R., 'All "dark patterns" are "hostile patterns": A hostility framework for understanding problematic digital interfaces', Ethics and Information Technology, Vol. 27, 2025, p. 57, https://link.springer.com/article/10.1007/s10676-025-09856-z.

UNESCO, Recommendation on the ethics of neurotechnology, 2025, https://www.unesco.org/en/legal-affairs/recommendation-ethics-neurotechnology.

UNESCO, 'Ethics of neurotechnology', UNESCO website, undated, accessed 30 June 2026, https://www.unesco.org/en/ethics-neurotech.

UN Human Rights Council Advisory Committee, Impact, opportunities and challenges of neurotechnology with regard to the promotion and protection of all human rights, United Nations, 2024, https://digitallibrary.un.org/record/4060417?v=pdf#files.

U.S. Senate Committee on Commerce, Science & Transportation, 'Sens. Cantwell, Schumer, Markey introduce legislation to shield Americans' brain data from exploitation', U.S. Senate Committee on Commerce, Science & Transportation website, 24 September 2025, accessed 30 June 2026, https://www.commerce.senate.gov/press/dem/release/sens-cantwell-schumer-markey-introduce-legislation-to-shield-americans-brain-data-from-exploitation/.

van den Hoven, J., Vermaas, P. E. and van de Poel, I. (Eds.), Handbook of Ethics, values, and technological design: Sources, theory, values and application domains. Springer, 2015.

Wascher, E., Heppner, H. and Hoffmann, S., 'Towards the measurement of event-related EEG activity in real-life working environments', International Journal of Psychophysiology, Vol. 91(1), 2014, pp. 3–9.

Winn, Z., 'Beacon Biosignals is mapping the brain during sleep', MIT News website, 1 May 2026, accessed 30 June 2026, https://news.mit.edu/2026/beacon-biosignals-maps-brain-during-sleep-0501.

World Health Organization (WHO), Policy brief on the ethics and governance of neurotechnology and AI, 2026.

Youngren, W.A., Horowitz, A.H., West Staples, V., Carr, M,, Stickgold, R. and Maes, P., ‘Targeted dream incubation and dream self-efficacy’, Sleep Adv, Vol. 6(2), 2025, zpaf013, https://academic.oup.com/sleepadvances/article/6/2/zpaf013/8063926.

Yuste, R., Goering, S., Arcas, B. A., Bi, G., Carmena, J. M. et al., ‘Four ethical priorities for neurotechnologies and AI’, Nature, Vol. 551(7679), 2017, pp. 59–163, https://www.nature.com/articles/551159a.

Zuboff, S., The age of surveillance capitalism: The fight for a human future at the new frontier of power, Public Affairs, 2019.

THE MEMBERS OF THE EGE

1. **Barbara PRAINSACK**
   Chair of the EGE
2. **Maria DO CÉU PATRÃO NEVES**
   Vice-Chair of the EGE
3. **Nils-Eric SAHLIN**
   Vice-Chair of the EGE
4. **Nikola BILLER-ANDORNO**
5. **Migle LAUKYTE**
6. **Paweł ŁUKÓW**
7. **Fruzsina MOLNÁR-GÁBOR**
8. **Thérèse MURPHY**
9. **Tamar SHARON**
10. **Marcel Jeroen VAN DEN HOVEN**
11. **Takis VIDALIS**

GETTING IN TOUCH WITH THE EU

In person

All over the European Union there are hundreds of Europe Direct centres. You can find the address of the centre nearest you online (european-union.europa.eu/contact-eu/meet-us_en).

On the phone or in writing

Europe Direct is a service that answers your questions about the European Union.
You can contact this service:

- by freephone: 00 800 6 7 8 9 10 11 (certain operators may charge for these calls),
- at the following standard number: +32 22999696,
- via the following form: european-union.europa.eu/contact-eu/write-us_en.

FINDING INFORMATION ABOUT THE EU

Online

Information about the European Union in all the official languages of the EU is available on the Europa website (european-union.europa.eu).

EU publications

You can view or order EU publications at op.europa.eu/en/publications. Multiple copies of free publications can be obtained by contacting Europe Direct or your local documentation centre (european-union.europa.eu/contact-eu/meet-us_en).

EU law and related documents

For access to legal information from the EU, including all EU law since 1951 in all the official language versions, go to EUR-Lex (eur-lex.europa.eu).

EU open data

The portal data.europa.eu provides access to open datasets from the EU institutions, bodies and agencies. These can be downloaded and reused for free, for both commercial and non-commercial purposes. The portal also provides access to a wealth of datasets from European countries.

This Statement by the European Group on Ethics in Science and New Technologies (EGE) focuses on the rise of advanced neuro-AI infrastructures—the integrated systems of hardware and software through which neurodata are collected, processed and reused to create, deploy and manage AI-powered neurotechnology applications.

Its central concern is how existing rights and ethical principles should be operationalised in relation to emerging neuro-AI infrastructures, including questions of European technological sovereignty, public-interest infrastructures and democratic accountability.

The Statement shifts the emphasis from neurodata privacy and isolated devices or applications to the collective impacts of neuro-AI infrastructures. It argues that we must move beyond a focus on data governance and recognise that the convergence of neurodata, neurotechnology and AI infrastructures requires a policy response that takes a systemic approach across regulatory domains.

Therefore, the EGE recommends: explicitly listing neurodata and neurodata-derived inferences as a distinct category of sensitive data, establishing specific governance requirements for Brain Foundation Models, Human (centric) Foundation Models, and neuro-AI infrastructures, establishing clear prohibitions where neuro-AI systems are used in consequential contexts, building European public-interest governance capacity for neuro-AI infrastructure, and conducting a targeted review of existing regulatory frameworks.

This Statement is published in the context of the Ethics Advice Mechanism (EthAM), which provides independent ethical advice and policy recommendations to European institutions on any topics where policymaking intersects with the development of science and technology.

Studies and reports

Publications Office
of the European Union