Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Relational Neurogovernance in Africa: Harambee, Botho, and Zambian Humanism as Foundations for Governing Neural Data and Brain–Computer Interfaces

Edward Musole

emusole@myune.edu.au

University of New England

Research Article

Keywords: neurotechnology governance, neural data, brain–computer interfaces, Ubuntu, relational ethics, African philosophy, neurorights, data protection

Posted Date: August 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10332787/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Governance frameworks for neurotechnology are emerging first in high-capacity jurisdictions whose regulatory instruments are built on an individualist conception of personhood. This article argues that African relational philosophies — Harambee, Botho, and Zambian Humanism — supply a normatively distinct and institutionally workable foundation for governing neural data and brain–computer interfaces (BCIs) in African states, addressing collective and relational harms that individual-consent-based models are not designed to reach. Rather than treating Western regulatory instruments as templates for direct transplantation, the article develops a method of contextualised transferability that separates specific technical mechanisms — tiered risk classification, third-party conformity assessment, and regulatory sandboxes — from the normative assumptions in which they originated, drawing comparatively on the European Union's Artificial Intelligence Act, Singapore's Model Artificial Intelligence Governance Framework, and Australia's sectoral model. These mechanisms are re-embedded within a three-stage African Relational Governance Framework (ARGF), anchored respectively in Harambee, Botho, and Zambian Humanism, with each stage linked to a specific protected value, institutional mechanism, and accountability metric. The framework is connected explicitly to existing continental instruments — the African Union Data Policy Framework and the Malabo Convention — so that it builds on, rather than duplicates, Africa's own regional governance architecture. The article concludes that relational neurogovernance is best understood not as a rejection of technical regulatory learning from elsewhere, but as a normative recalibration of it.

1. Introduction

1.1 What this article is about

Neurotechnology — devices and systems that record, decode, monitor, or modulate nervous-system activity — is no longer confined to clinical research. Consumer electroencephalography (EEG) headsets, implantable brain–computer interfaces (BCIs) for motor and communication restoration, closed-loop neurostimulation devices, and neural-data-driven consumer analytics are all in active development or deployment [1, 2]. The Organisation for Economic Co-operation and Development (OECD) formally recognised the governance urgency of this trajectory in December 2019, when it adopted the Recommendation on Responsible Innovation in Neurotechnology, the first international instrument dedicated specifically to this technology class [1]. Chile followed in October 2021 by amending Article 19 of its Constitution to protect mental integrity against the effects of neurotechnology, becoming the first state to entrench "neurorights" in constitutional law [3]. Chile's move has not remained isolated: legislators in Brazil, Mexico, Argentina, and Colombia have since advanced their own neurorights bills and charters, though a recent peer-reviewed assessment of a regional model law developed by the Latin American Parliament argues that this proliferation has reproduced rather than resolved the conceptual ambiguities of the underlying neurorights framework, including overly broad drafting and weak scientific grounding [4]. This is a useful caution for the analysis that follows: the existence of comparator legislation elsewhere is not by itself evidence that the underlying legal concepts are settled. These are

not speculative future problems; they are live governance questions being actively settled now, largely by institutions outside Africa.

The urgency of this trajectory is not confined to policy instruments. In 2018, the Global Neuroethics Summit Delegates called on each of the world's major national brain-research initiatives to build neuroethics review into their governance structures from the outset rather than retrofitting it later [5]. The following year, an international group convened as the Morningside Group proposed four ethical priorities — privacy and consent, agency and identity, augmentation, and bias — specifically for neurotechnology and artificial intelligence, on the view that existing ethical guidance for either technology alone was insufficient for systems combining both [6]. None of the seven brain initiatives that prompted the 2018 call originates in Africa, and none of the subsequent instruments discussed in this article was drafted with African legal or institutional conditions as a reference case. The pattern is consistent: the ethical and regulatory vocabulary for neurotechnology is being settled elsewhere, in advance of the deployment decisions African states will eventually have to make.

This creates what we term a sovereignty paradox: by the time neurotechnologies reach widespread deployment in African markets, the governing norms may already have solidified elsewhere, leaving domestic regulators to adopt frameworks whose underlying assumptions about personhood, consent, and harm were never built with African social and institutional conditions in mind [7]. Eke's analysis of neurotechnology governance in Africa argues that the continent should draw lessons from prior experience with artificial intelligence (AI) governance precisely in order to avoid repeating that pattern — designing anticipatory, contextually grounded rules before deployment outpaces regulatory capacity [2].

This article asks a narrower and more operational question than whether Africa should import Western technology law: given that some technical regulatory mechanisms — risk tiering, conformity assessment, sandboxes — are genuine engineering solutions to real oversight problems, how can African jurisdictions adopt what is technically useful in them while grounding the normative commitments — what counts as harm, who must be consulted, what consent means — in ethical traditions with actual purchase in the societies being regulated?

1.2 Key terms

Two terms are used throughout and are defined here rather than left implicit.

Neurotechnology, for the purposes of this article, refers to technologies that record, decode, monitor, or modulate neural activity, spanning a spectrum from non-invasive consumer EEG and neurofeedback devices, through diagnostic and therapeutic tools such as deep brain stimulation, to invasive BCIs for communication and motor restoration [1]. Neurotechnology is treated as a subset of, but analytically distinct from, artificial intelligence: many neurotechnologies incorporate AI-based signal decoding, but the ethical and legal stakes specific to neurotechnology — the biological source of the data and its potential to reveal identity- and cognition-linked information — are not fully captured by AI governance instruments alone [8].

Global North and Global South are used descriptively, not as fixed geographic categories, to denote asymmetries in regulatory capacity, standard-setting power, and technology-exporting versus technology-importing positions in the current neurotechnology governance landscape [7]. The terms carry no claim about the internal homogeneity of either grouping; African states themselves differ substantially in institutional capacity, and this is addressed directly in Section 3.

1.3 Contribution and scope

This article makes three contributions. First, it identifies a specific normative gap in individual-consent-centred neurodata governance — its limited capacity to address relational and collective harms — and argues that this gap is addressed more directly by relational African ethical traditions than by supplementary technical fixes within the individualist paradigm [8, 9]. Second, it develops a method of contextualised transferability, adapting Rose's concept of lesson-drawing [10] and Pistor's account of legal transplantation [11], to separate technical governance mechanisms from the normative architecture in which they were designed, applying this method across three jurisdictions with different institutional logics — the European Union, Singapore, and Australia. Third, it operationalises the resulting African Relational Governance Framework (ARGF) as a three-stage model in which each stage specifies a protected value, a concrete institutional mechanism, a responsible actor, and a review metric, connected explicitly to the African Union's own data governance instruments.

The analysis is illustrative rather than continent-representative. It draws on Harambee (Kenya), Botho (Botswana), and Zambian Humanism as three articulated relational philosophies with a documented history of translation into public policy [12–14], while acknowledging — as Section 3 discusses — that they do not exhaust Africa's philosophical diversity.

2. Why Individual-Consent Models Are Incompletely Specified for Neural Data

Contemporary data protection instruments are built around individual consent as the principal legitimating mechanism for data processing [15]. This model has been effective for a wide range of privacy and consumer-protection problems. It is, however, poorly matched to a structural feature of neural data: neural signals frequently encode information that extends beyond the individual from whom they were recorded, including inferences bearing on family members, linguistic and cultural group membership, and shared cognitive or behavioural traits [8]. Ienca and Andorno's account of emerging neurorights was itself motivated by the observation that existing human-rights and data-protection instruments do not adequately address these group-level and identity-linked dimensions of neural data [8].

This is not a claim that individual consent is unimportant, or that EU- or Australian-style instruments are poorly designed for the problems they were built to solve. It is a claim that a governance framework organised exclusively around individual consent will systematically under-govern a category of harm — communal exposure, cultural-group inference, relational privacy loss — that is especially salient for

neural data. Frameworks grounded in relational conceptions of personhood treat the person as constituted through relationships from the outset, which gives them a more natural conceptual vocabulary for such harms without needing to bolt on ad hoc exceptions to an individualist baseline [9, 16].

This critique of individualism is not unique to African relational philosophy, and the article does not present it as such. Mainstream bioethics and feminist philosophy have developed a substantial parallel literature on relational autonomy, arguing that autonomous agency itself is socially constituted through relationships of care, dependency, and mutual recognition rather than exercised by a pre-social, self-sufficient chooser [17, 18]. That literature reaches a similar diagnosis of individual consent's limits by a different route, and its existence matters for this article's argument: it shows that the case for a relational starting point in data governance is not a claim that depends on accepting a specifically African metaphysics, even though this article develops that case through African relational philosophy in particular, for the institutional-design reasons set out in Section 3.

A parallel scholarly response to this gap has developed within neurorights and neurolaw scholarship, which has proposed a series of new individual rights specific to neurotechnology, including rights to mental privacy, personal identity, free will, and fair access, and the broader concept of cognitive liberty, or mental self-determination, as a legal foundation for regulating interventions into the mind [6, 8, 19]. This literature is a valuable complement to, but not a substitute for, the argument advanced here. Cognitive liberty and the associated neurorights proposals are framed, deliberately, as individual entitlements against interference with one’s own mind; they extend the individualist paradigm to cover a new object (the mind) rather than replacing its unit of analysis (the individual). An international consensus statement by Goering, Klein, and colleagues makes a similar point from within neuroethics itself: it recommends new individual-facing protections for identity, agency, privacy, and bias, but does not propose a mechanism for registering harms that fall on a community rather than an identified individual [20]. Such proposals therefore strengthen the case that current instruments under-protect neurotechnology-specific interests, but they do not, on their own terms, supply the relational or collective register this article argues is separately needed.

3. African Relational Philosophy as a Normative Foundation

Ubuntu — captured in the maxim "I am because we are" — is the most widely discussed instance of a broader family of relational ontologies found across African philosophical traditions, in which personhood is achieved through, rather than prior to, social relationship and moral participation [9, 16]. Metz's systematic treatment of this tradition develops it into a comprehensive moral theory grounded in communal relationship rather than individual utility or autonomy, with application across bioethics, research ethics, and other applied domains [16]. Mhlambi's account applies the same relational ontology specifically to the governance of automated decision-making systems, arguing that many of the harms such systems produce are best understood as violations of relational personhood rather than as isolated infringements of individual rights [9].

The case for relational ethics made in this article is not simply that AI governance in general benefits from a relational lens, which is a claim already well established in the decolonial-AI literature discussed below; it is a narrower claim about why neural data is a distinctively strong case for that lens, stronger than most other data categories to which relational critiques of individualism have been applied. Three features of neural data make this so. First, neural data is constitutively, not merely statistically, relational: because brain activity encodes linguistic processing, family-linked genetic and health markers, and culturally shared cognitive patterns, a single person's neural record discloses information about their kin and community as a structural feature of what the data is, not as an incidental correlation that better anonymisation could remove. Second, neural data implicates identity and agency in a way most consumer data does not: an intervention that decodes or modulates neural activity acts on the very faculties by which a person would otherwise object to, consent to, or renegotiate the intervention, which is not true of, for example, a retailer inferring purchasing preferences. Third, the harms at stake are frequently irreversible at the level of the person rather than merely costly to remedy: a misused shopping profile can be deleted, but an identity-altering neurostimulation effect or a discriminatory inference drawn from decoded neural patterns cannot always be undone by after-the-fact correction. Taken together, these features mean that neural data is not simply one more instance where a relational critique of individualism happens to apply; it is a category where the individualist starting point is least defensible on its own terms, because the boundary between "my data" and "our data" is harder to draw coherently for neural information than for almost any other data type currently subject to protection frameworks.

This article works with three specific articulations of this broader relational orientation:

1. Zambian Humanism, articulated by Kenneth Kaunda as a national governing philosophy grounded in the interdependence of persons within community [13];
2. Botho, the Tswana/Sotho principle that personhood is achieved through respectful and dignified relationship [14];
3. Harambee, the Kenyan ethic of "pulling together" that treats collective responsibility and mutual aid as foundational to social organisation [12].

These three were selected because each has a documented history of translation from philosophical articulation into public policy, which matters for a framework whose purpose is institutional design rather than purely conceptual analysis. They are not offered as jointly exhaustive of African ethical thought, nor as statistically representative of a continent comprising 54 states and a far larger number of ethnic and linguistic traditions. The claim made here is narrower: that these three traditions, precisely because they have already been operationalised into governance discourse in their respective countries, illustrate how a relational ontology can be translated into institutional design principles — for consent, for harm, and for accountability — that a purely individualist framework does not generate on its own terms.

This positioning is a deliberate narrowing relative to earlier work in this area, which used a similar three-country selection but advanced a stronger, continent-wide representativeness claim. The claim here is illustrative and methodological: relational ontologies of this kind recur across many African traditions, and the three used here are a tractable entry point into operationalisation, not a survey of that diversity.

The recurrence of this relational ontology beyond the three countries discussed here is not merely asserted; it is the subject of a long-running debate within African philosophy itself, most closely associated with Menkiti and Gyekye. Menkiti's influential account argues that personhood in African traditional thought is conferred by the community and is something an individual can, in principle, fail fully to attain, making the community's reality prior to that of the individual life history [21]. Gyekye's response defends a more moderate communitarianism [22], in which the person is still constituted through social relationship but retains a residual capacity for individual evaluation and dissent that Menkiti's stronger formulation leaves too little room for. This article's operationalisation of the ARGF in Section 5 is deliberately closer to Gyekye's moderate position than to Menkiti's stronger one: the framework treats communal assessment and individual consent as independent, non-substitutable gates precisely so that a community-level finding can never be used to override a specific individual's informed refusal (Section 5.5). This is a substantive philosophical choice, not merely an implementation detail, and it is one place where this article takes a position within African philosophy rather than simply reporting on it.

This project also sits within a broader decolonial-technology literature that has, over the past several years, examined how AI development and deployment can reproduce colonial-era extraction patterns and epistemic hierarchies unless deliberately redesigned [23]. Birhane's account of algorithmic injustice makes the connection to relational African ethics explicit, arguing that predictive systems built on an individualist, atomistic conception of the person will systematically misdescribe socially embedded harms, and proposing relational ethics — including but not limited to Ubuntu — as a corrective lens for data science more broadly [24]. A recent empirical assessment of AI governance across Sub-Saharan Africa similarly finds that most existing national AI strategies on the continent remain modelled closely on EU or OECD templates, with limited adaptation to local governance traditions, which is the specific pattern this article is trying to interrupt for the neurotechnology sub-case [25]. This literature is concerned primarily with critique and with the demand for epistemic justice in whose knowledge counts in designing governance; this article shares that concern but is oriented toward a narrower, constructive task — showing one concrete way relational African philosophy can be translated into licensing criteria, review metrics, and institutional mandates for a specific technology class, rather than offering a general theory of technological decolonisation.

4. Comparative Governance Architectures: Three Technical Archetypes

Three jurisdictions are examined here as sources of technical governance mechanisms — not as normative benchmarks — because together they represent three structurally different approaches to

technology regulation: comprehensive codified risk-tiering (European Union), principles-based agile guidance (Singapore), and sectoral incremental adaptation (Australia). This comparison is conceptual and doctrinal rather than an empirical content analysis of the underlying regulatory texts: it draws on secondary legal-scholarship characterisations of how each instrument operates in practice, not on original coding of statutory or guidance language, and the article's contribution should be read accordingly as comparative-institutional rather than empirical. Singapore, rather than a South Asian jurisdiction, is used here as the principles-based comparator because it is the most extensively documented exemplar of agile, non-binding AI governance in the literature this article engages with directly; a fuller comparative study could usefully extend this method to emerging South Asian frameworks as an additional test of its transferability.

4.1 The European Union: risk-tiered codification

The EU's Artificial Intelligence Act, which entered into force in August 2024, classifies AI systems into unacceptable-risk, high-risk, limited-risk, and minimal-risk tiers, with high-risk systems — a category that would capture most clinical and many consumer neurotechnologies — subject to conformity assessment, risk-management obligations, and human oversight requirements [26]. The Act's use of accredited third-party conformity-assessment bodies distributes technical evaluation capacity beyond the state itself, which is potentially valuable for jurisdictions without large in-house technical regulatory bureaucracies.

From a relational perspective, the Act's limitation is definitional rather than technical: its risk categories are specified in terms of individual safety and fundamental rights, without an explicit mechanism for communal consultation or for registering harms that are collective or relational in character [8, 26]. This is a gap in scope, not a design flaw relative to the Act's own stated purpose.

4.2 Singapore: principles-based, agile governance

Singapore's Model Artificial Intelligence Governance Framework, published by the Personal Data Protection Commission and the Infocomm Media Development Authority, takes a materially different approach: voluntary, principles-based guidance organised around internal governance structures, human-centric design, and operations management, rather than binding ex ante classification [27]. Singapore's model is instructive for two reasons distinct from the EU's: it demonstrates that a credible governance signal can be built without comprehensive binding legislation, and it demonstrates iterative institutional learning — the framework has been revised repeatedly since 2019 as the underlying technology changed [27]. Its limitation for direct transfer is that its legitimacy rests substantially on Singapore's unusually high institutional trust and regulatory capacity, conditions that cannot be assumed uniformly across African jurisdictions.

4.3 Australia: sectoral incrementalism

Australia does not have a single comprehensive neurotechnology or AI statute. Instead, existing sectoral regulators extend their mandates through interpretation and guidance — therapeutic neurotechnologies

fall within the Therapeutic Goods Administration's medical-device framework, while financial-sector regulators have used regulatory sandboxes to permit supervised experimentation. This model's advantage is that it avoids duplicating institutional capacity; a "neuro-unit" can be layered into an existing health or data-protection regulator rather than requiring a new bureaucracy. Its disadvantage, from a relational perspective, is fragmentation: without a unifying normative statement, sectoral extension risks producing inconsistent, ad hoc rules with no forum for community input.

4.4 What is transferable, and what is not

None of the three models is normatively transferable as a whole; each rests on institutional and cultural preconditions — a constitutionalised rights tradition in the EU's case, exceptionally high regulatory trust in Singapore's, and a mature independent judiciary in Australia's — that do not travel automatically. What is transferable is a specific, bounded set of technical mechanisms: tiered risk classification as a triage device for allocating scarce regulatory attention; third-party conformity assessment as a way of extending technical capacity without a large permanent bureaucracy; and regulatory sandboxes as a way of generating regulatory knowledge about a new technology before committing to permanent rules. Table 1 sets out this extraction explicitly, including the relational critique of each mechanism's default framing and the recalibration required to close that gap.

Table 1
Comparative Governance Architectures: Transferable Mechanisms and Relational Recalibration

Governance Model	Transferable Technical Mechanism	Institutional Precondition	Relational Critique	Recalibration for ARGF
European Union (AI Act)	Tiered risk classification; third-party conformity assessment	Constitutionalised rights tradition; large technical bureaucracy	Risk defined by individual safety/rights only; no communal-harm category	Add collective-impact criteria to risk tiers; require Community Impact and Dignity Assessment for high-risk systems
Singapore (Model AI Governance Framework)	Principles-based agile guidance; iterative revision cycle	Very high institutional trust and regulatory capacity	Voluntary compliance model presumes trust conditions many African states cannot assume	Retain iterative-revision logic; make community consultation a mandatory (not voluntary) element
Australia (sectoral incrementalism)	Regulatory sandboxes; institutional layering within existing agencies	Mature independent judiciary; established sectoral regulators	Fragmentation risk; no unifying normative statement across sectors	Layer a neuro-unit into an existing data-protection or health regulator, under the unifying ARGF statement

4.5 The regional anchor: existing African instruments

Any relational governance model proposed for African jurisdictions must be positioned relative to instruments that already exist at the continental level. The African Union's Data Policy Framework, endorsed by the AU Executive Council in February 2022 and published in July 2022, sets out continental guidance on data governance, cross-border data flow, and the balance between innovation and protection in Africa's digital economy [28]. The African Union Convention on Cyber Security and Personal Data Protection (the Malabo Convention), adopted in 2014, entered into force in June 2023 after securing its fifteenth ratification, and is the only binding continent-wide treaty addressing data protection and cybersecurity in Africa [29]. Both instruments are silent on neurotechnology specifically, which is the gap the framework proposed in Section 5 is intended to fill — not as a competing structure, but as a sector-specific elaboration that a state already committed to the Malabo Convention or the AU Data Policy Framework could adopt as an extension of existing commitments.

5. The African Relational Governance Framework (ARGF)

The ARGF is a three-stage maturity model. Each stage is anchored in one of the three relational philosophies discussed in Section 3, and each stage specifies four elements: the value the stage is designed to protect, the institutional mechanism that operationalises that protection, the actor responsible for the mechanism, and a metric against which its performance can be reviewed. Table 2 summarises this structure across all three stages, setting out the protected value, mechanism, actor, and review metric for each in turn.

5.1 Stage 1 — Foundational Collective Protection (Harambee)

Protected value

collective safety of digital and neural infrastructure, treated as a shared responsibility rather than solely an individual entitlement, consistent with Harambee's foregrounding of collective responsibility and mutual aid [12].

Mechanism

baseline data-protection legislation that explicitly recognises collective and communal data harms as a ground for regulatory action; an independent data protection authority with an express mandate covering both individual and collective harms; alignment with the cybersecurity obligations already specified in the Malabo Convention rather than duplication of them [29].

Actor

national data protection authority, operating within the AU Data Policy Framework's coordination architecture [28].

Review metric

documented incidence and remediation time for communal (not only individual) data-breach events; ratification and domestic implementation status of the Malabo Convention.

5.2 Stage 2 — Risk-Tiered Relational Oversight (Botho)

Protected value

relational dignity — the idea, drawn from Botho, that a person is constituted through respectful relationship, which requires that neurotechnology licensing assess social and communal impact and not only individual technical safety [14].

Mechanism

classification of neural data as a high-sensitivity category within the national data-protection regime; a Community Impact and Dignity Assessment as a licensing precondition for implantable or clinically significant neurotechnologies, run alongside — not instead of — individual informed-consent procedures; an Algorithmic Impact Assessment evaluating group-level bias in neural-signal decoding models.

Actor

sectoral regulator (health products authority or equivalent), applying the risk-tiering logic transferred from the EU AI Act (Section 4.1) but with assessment criteria widened to include the collective-harm categories identified in Section 2.

Review metric

proportion of high-risk neurotechnology licences accompanied by a completed Community Impact and Dignity Assessment; documented instances in which such an assessment altered licensing conditions.

5.3 Stage 3 — Structural Interdependence (Zambian Humanism)

Protected value

long-term interdependence between neurotechnology users, providers, and the wider community, consistent with Zambian Humanism's emphasis on interdependence as the ground of human dignity [13].

Mechanism

legacy-support escrow arrangements requiring BCI and implant providers to fund continuity-of-care provisions in the event of company failure; open, non-proprietary neural-signal interoperability standards to prevent vendor lock-in for implant recipients; participation in a regional maintenance consortium, building on the coordination role the AU Data Policy Framework assigns to Regional Economic Communities [28].

Actor

sector regulator in coordination with regional economic communities and, where relevant, the African Union.

Review metric

number of implant recipients covered by continuity-of-care escrow; existence and membership of a functioning regional maintenance consortium.

Table 2
The African Relational Governance Framework (ARGF)

Stage	Protected Value	Mechanism	Actor	Review Metric
1. Foundational Collective Protection (Harambee)	Collective digital and neural safety	Baseline data-protection law recognising collective harms; DPA with collective-harm mandate; alignment with Malabo Convention	National data protection authority	Communal breach incidence/remediation time; Malabo Convention implementation status
2. Risk-Tiered Relational Oversight (Botho)	Relational dignity	Neural data classified high-sensitivity; Community Impact and Dignity Assessment (alongside individual consent); Algorithmic Impact Assessment for group-level bias	Sectoral (health products) regulator	Proportion of high-risk licences with completed assessment; documented licence-condition changes
3. Structural Interdependence (Zambian Humanism)	Long-term interdependence of users, providers, community	Legacy-support escrow funds; open interoperability standards; regional maintenance consortium	Sector regulator with Regional Economic Communities / AU	Implant recipients covered by escrow; consortium existence/membership

5.4 Anchoring the ARGF in existing domestic law

The ARGF is not proposed against a blank legislative slate in Kenya, Zambia, or Botswana. Each of the three countries whose relational philosophy anchors an ARGF stage already has an operative data-protection statute, and Stage 1 in particular is best read as a proposal to extend, rather than replace, these existing instruments.

Kenya's Data Protection Act, 2019 (No. 24 of 2019) establishes an Office of the Data Protection Commissioner with a broad mandate over the processing of personal data, modelled substantially on the EU's General Data Protection Regulation [30]. Zambia's Data Protection Act, 2021 (No. 3 of 2021) established an analogous data protection authority and regime, enacted alongside a companion Cyber Security and Cyber Crimes Act in the same legislative year [31]. Botswana's Data Protection Act, 2018 (No. 32 of 2018), substantially revised by the Data Protection Act, 2024 (No. 18 of 2024), established the Information and Data Protection Commission [32]. None of these three statutes contains a category of

neural or neurotechnology-specific data, and none makes explicit provision for communal or collective data harms distinct from harms to an identified individual.

This is precisely the gap Stage 1 of the ARGF is designed to close — not by displacing these statutes but by amending them: adding neural data as an explicitly recognised category of sensitive personal data, in the way several jurisdictions already treat health and biometric data, and adding a collective-harm ground for regulatory action that operates alongside, not instead of, the individual-harm grounds these Acts already recognise. Because all three jurisdictions already have a functioning data protection authority, Stage 1 does not require new institutional capacity, only an expanded mandate for an existing institution — which is precisely the kind of institutional layering identified as a strength of the Australian model in Section 4.3, applied here to instruments that already exist in the very countries the ARGF's founding philosophies come from.

5.5 Addressing the limits of a three-philosophy framework

Three objections deserve direct engagement. First, a framework built on three national philosophies cannot claim to represent 54 states; this article's response is to narrow the claim to illustration and method rather than representativeness (Section 3). This objection is sharpened by Mboti's critique of what he calls the "greatest harmony" account of Ubuntu, which argues that much of the secondary literature invoking Ubuntu treats African communitarianism as a uniform, harmony-obsessed essence, flattening real disagreement within African philosophy about the relationship between individual and community [33]. The response adopted here is to keep the claim at the level Mboti's critique leaves open: that three specific, nationally articulated philosophies generate specific institutional design principles, without asserting that they exhaust or harmonise African ethical thought. Second, communal consultation mechanisms of the kind proposed in Stage 2 create a genuine risk of privileging the loudest or most institutionally connected community voices over dissenting individuals, particularly minorities within a community; the Community Impact and Dignity Assessment is therefore specified to run alongside, not in place of, individual informed consent. Concretely, this means a negative Community Impact and Dignity Assessment can block a provider's licence to operate or market a technology in a given setting, and can attach binding conditions (community liaison, benefit-sharing, enhanced monitoring) to a licence that is granted, but it cannot be used to override a specific individual's informed refusal, and a positive assessment cannot be used to pressure an individual into consenting. Symmetrically, an individual's consent does not substitute for a Community Impact and Dignity Assessment where one is required; the two operate as independent, non-substitutable gates rather than as inputs to a single combined score, precisely so that neither can be trimmed to unblock the other administratively. Any residual conflict, such as a community body's objection to a technology that a fully informed individual still wishes to use for herself alone, is resolved in favour of the individual once the assessment has confirmed the effects are genuinely confined to her; the community-level gate exists to catch harms that spill beyond the individual, not to license paternalism over purely self-regarding choices. Third, relational governance mechanisms could in principle be captured by authoritarian governments seeking to justify restrictions on individual neurotechnology access in the name of

"communal welfare"; guarding against this requires that any communal consultation body include independent civil-society representation and be subject to judicial or quasi-judicial review — a precondition for legitimate use of the framework, not an optional addition.

6. A Worked Example: Licensing an Implantable BCI Trial Under the ARGF

The mechanisms proposed in Section 5 are easiest to assess against a concrete scenario. Consider a foreign medical-device company seeking to run a clinical trial of an implantable BCI for motor restoration at a hospital in Zambia, recruiting patients with severe spinal-cord injury, with neural signal data processed on cloud infrastructure hosted outside Zambia.

Under the status quo described in Section 4.5, this trial would be reviewed under Zambia's Data Protection Act, 2021 as an instance of sensitive personal data processing, requiring individual informed consent from each participant and, depending on classification, a data protection impact assessment focused on the individual participants named in the trial protocol [31]. Nothing in that review would require the sponsor to consider whether decoded neural signals might carry inferential information about participants' extended families, nor whether the trial's data-export arrangements leave Zambia dependent on a single foreign infrastructure provider for the continued function of devices already implanted in Zambian citizens.

Under Stage 1 of the ARGF, the trial would additionally trigger the amended collective-harm ground for regulatory action proposed in Section 5.1: the Data Protection Commissioner would have express authority to require the sponsor to disclose whether the decoding pipeline could plausibly reveal information about participants' relatives, given that motor-BCI decoding models are frequently trained on and validated against family-linked genetic and neurophysiological reference data. Under Stage 2, the trial would require a Community Impact and Dignity Assessment before licensing, run by the sectoral regulator alongside the standard individual-consent process, examining whether recruitment, monitoring, or device-maintenance arrangements risk stigmatising participants within their communities, and an Algorithmic Impact Assessment examining whether the sponsor's decoding model has been validated on populations resembling the Zambian participant pool or only on demographically different training data. Under Stage 3, the licence would be conditioned on the sponsor's participation in a legacy-support escrow arrangement funding continued device maintenance if the sponsor exits the market, and on the use of interoperable, non-proprietary signal formats so that a participant's implant remains serviceable by another provider if the original sponsor withdraws.

What this adds beyond the status quo is not a duplicate layer of individual consent, but three things ordinary informed consent does not reach on its own: an explicit check for family- and community-linked inferential harms in the decoding pipeline itself, a formal assessment of whether the trial's algorithmic tools were built for a population resembling the one being recruited, and a durable continuity guarantee for implant recipients that does not depend on the sponsor's continued commercial presence in Zambia.

Equally, the worked example shows the limit of the framework: it does not resolve whether the trial should proceed at all, nor does it substitute for the sponsor's own scientific and safety review under the Therapeutic Goods-style clinical trial pathway a Zambian regulator would separately require. It adds specific, checkable requirements to an approval process that would otherwise be silent on collective and structural risk.

7. Discussion: Neuro-Governance as Institutional Capital

A jurisdiction that invests early in credible, contextually legitimate neurotechnology governance is better positioned to negotiate licensing and data-localisation terms from a position of institutional credibility rather than institutional absence [11]. This is not a claim that governance capacity substitutes for other developmental priorities; it is a narrower claim that the sequencing and normative grounding of early governance choices in a nascent technology market have effects on subsequent bargaining position that are more durable than accounts treating regulation purely as a cost to be minimised tend to assume [11].

8. Limitations and Directions for Future Research

Three limitations bound the claims made here and set an agenda for further work. First, the framework is normative and institutional-design-oriented rather than empirical: it has not been tested against actual licensing decisions, and the review metrics proposed in Section 5 are proposed as a starting design rather than validated instruments. A natural next step is a small-sample study of how a Community Impact and Dignity Assessment would apply to an actual pending neurotechnology deployment in one of the three anchor jurisdictions, to test whether the assessment generates decision-relevant information beyond what existing individual consent processes already capture.

Second, the ARGF's civil-society and judicial-review safeguards (Section 5.5) are stated as design requirements rather than demonstrated as achievable in every institutional context; their feasibility will vary with the independence of the judiciary and the strength of civil society in a given jurisdiction, and the framework offers no mechanism for what should happen where those preconditions are absent beyond noting that the framework should not be considered legitimately implemented without them.

Third, this article has not addressed the question of how the ARGF would interact with cross-border neural data flows, which is likely to be the more common practical scenario than purely domestic processing, given that most neurotechnology hardware and cloud-based decoding infrastructure is currently produced and hosted outside Africa. Extending the framework to cross-border transfer conditions, potentially building on the Malabo Convention's existing (non-neurotechnology-specific) transfer provisions [29], is a priority for further work.

9. Conclusion

Neurotechnology governance is being actively constructed now, mostly by institutions outside Africa, and the normative assumptions embedded in that construction will not automatically fit African social

and institutional contexts if adopted wholesale. This article has argued for a specific alternative: separate the technical governance mechanisms that have genuine engineering value from the individualist normative architecture in which they were designed, and re-embed them within a relational ethical framework — illustrated here through Harambee, Botho, and Zambian Humanism — capable of registering the collective and relational harms neural data specifically raises. The resulting African Relational Governance Framework is offered as one operationalisation of that method, positioned alongside existing continental instruments and explicitly bounded in its representativeness claims.

Declarations

Conflicts of interest:

The author declares no conflict of interest.

Funding:

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contribution

E.M prepared and wrote the entire manuscript.

Acknowledgement

none

Data availability:

This is a conceptual article; no datasets were generated or analysed.

References

1. OECD (Organisation for Economic Co-operation and Development). 2019. *Recommendation of the Council on Responsible Innovation in Neurotechnology.* OECD/LEGAL/0457. Paris: OECD.
2. Eke, Damian. 2024. Ethics and governance of neurotechnology in Africa: Lessons from AI. *JMIR Neurotechnology* 3: e56665. https://doi.org/10.2196/56665.
3. Chile. 2021. Ley N° 21.383, que modifica el artículo 19, numeral 1°, de la Constitución Política de la República. *Diario Oficial de la República de Chile*.

4. Borbón, Diego. 2025. What a NeuroRights legislation should not look like: The case of the Latin American Parliament. *Frontiers in Neuroscience* 18: 1514338.
5. Global Neuroethics Summit Delegates, Karen S. Rommelfanger, Sung-Jin Jeong, Arisa Ema, et al. 2018. Neuroethics questions to guide ethical research in the international brain initiatives. *Neuron* 100: 19–36. https://doi.org/10.1016/j.neuron.2018.09.021.
6. Yuste, Rafael, Sara Goering, Blaise Agüera y Arcas, et al. 2017. Four ethical priorities for neurotechnologies and AI. *Nature* 551: 159–163. https://doi.org/10.1038/551159a.
7. Couldry, Nick, and Ulises A. Mejias. 2019. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford: Stanford University Press.
8. Ienca, Marcello, and Roberto Andorno. 2017. Towards new human rights in the age of neuroscience and neurotechnology. *Life Sciences, Society and Policy* 13: 5. https://doi.org/10.1186/s40504-017-0050-1.
9. Mhlambi, Sabelo. 2020. *From Rationality to Relationality: Ubuntu as an Ethical and Human Rights Framework for Artificial Intelligence Governance.* Carr Center Discussion Paper Series 2020-009. Cambridge, MA: Harvard Kennedy School.
10. Rose, Richard. 1993. *Lesson-Drawing in Public Policy: A Guide to Learning Across Time and Space*. Chatham, NJ: Chatham House Publishers.
11. Pistor, Katharina. 2019. *The Code of Capital: How the Law Creates Wealth and Inequality*. Princeton: Princeton University Press.
12. Ngau, Peter M. 1987. Tensions in empowerment: The experience of the 'Harambee' (self-help) movement in Kenya. *Economic Development and Cultural Change* 35: 523–538.
13. Kaunda, Kenneth D. 1967. *Humanism in Zambia and a Guide to Its Implementation, Part I*. Lusaka: Zambia Information Services.
14. Dolamo, Ramathate. 2013. Botho/Ubuntu: The heart of African ethics. *Scriptura* 112 (1): 1–10. https://doi.org/10.7833/112-0-78.
15. Bennett, Colin J., and Charles D. Raab. 2006. *The Governance of Privacy: Policy Instruments in Global Perspective*, 2nd ed. Cambridge, MA: MIT Press.
16. Metz, Thaddeus. 2022. *A Relational Moral Theory: African Ethics in and Beyond the Continent*. Oxford: Oxford University Press.
17. Mackenzie, Catriona, and Natalie Stoljar, eds. 2000. *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*. New York: Oxford University Press.
18. Held, Virginia. 2006. *The Ethics of Care: Personal, Political, and Global*. New York: Oxford University Press.
19. Bublitz, Jan Christoph. 2013. My mind is mine!? Cognitive liberty as a legal concept. In *Cognitive Enhancement: Trends in Augmentation of Human Performance*, Vol. 1, eds. Elisabeth Hildt and Andreas G. Franke, 233–264. Dordrecht: Springer.

20. Goering, Sara, Eran Klein, Laura Specker Sullivan, et al. 2021. Recommendations for responsible development and application of neurotechnologies. *Neuroethics* 14: 365–386. https://doi.org/10.1007/s12152-021-09468-6.
21. Menkiti, Ifeanyi A. 1984. Person and community in African traditional thought. In *African Philosophy: An Introduction*, 3rd ed., ed. Richard A. Wright, 171–181. Lanham: University Press of America.
22. Gyekye, Kwame. 1997. *Tradition and Modernity: Philosophical Reflections on the African Experience*. New York: Oxford University Press.
23. Mohamed, Shakir, Marie-Therese Png, and William Isaac. 2020. Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology* 33 (4): 659–684. https://doi.org/10.1007/s13347-020-00405-8.
24. Birhane, Abeba. 2021. Algorithmic injustice: A relational ethics approach. *Patterns* 2: 100205. https://doi.org/10.1016/j.patter.2021.100205.
25. Ayana, Gelan, Kokeb Dese, Hundessa Daba Nemomssa, et al. 2024. Decolonizing global AI governance: Assessment of the state of decolonized AI governance in Sub-Saharan Africa. *Royal Society Open Science* 11: 231994. https://doi.org/10.1098/rsos.231994.
26. European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union* L 2024/1689, 12 July 2024.
27. Personal Data Protection Commission Singapore, and Infocomm Media Development Authority. 2020. *Model Artificial Intelligence Governance Framework*, 2nd ed. Singapore: PDPC/IMDA.
28. African Union. 2022. *AU Data Policy Framework*. Addis Ababa: African Union.
29. African Union. 2014. *African Union Convention on Cyber Security and Personal Data Protection (Malabo Convention).* Adopted 27 June 2014; entered into force 8 June 2023.
30. Kenya. 2019. Data Protection Act, 2019 (No. 24 of 2019).
31. Zambia. 2021. Data Protection Act, 2021 (No. 3 of 2021).
32. Botswana. 2024. Data Protection Act, 2024 (No. 18 of 2024), substantially revising the Data Protection Act, 2018 (No. 32 of 2018).
33. Mboti, Nyasha. 2015. May the real ubuntu please stand up? *Journal of Media Ethics* 30: 125–147. https://doi.org/10.1080/23736992.2015.1020380.