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.

Beyond Neuro-Extractivism: Building a Neural Chilimba Framework for Culturally Grounded Neuro-Education Governance in Zambia

Edward Musole

emusole@myune.edu.au

University of New England

Samson Kantini

Zambia National Commission for UNESCO

Research Article

Keywords: neuro-extractivism, Ubuntu, Neural Chilimba, educational neurotechnology, data sovereignty, Zambia, LMIC governance

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9385556/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

Educational neurotechnology — portable electroencephalography (EEG) headsets, attention-tracking wearables, and cognitive-load monitors — is becoming commercially accessible in Global South markets with almost no neuro-specific regulatory infrastructure. Left ungoverned, these conditions create structural risk of neuro-extractivism: the asymmetric and largely irreversible flow of children's cognitive data from low- and middle-income country (LMIC) educational settings to technology corporations operating in regulatory vacuums. Existing neuroethics frameworks, grounded in individual consent, liberal privacy rights, and Western regulatory models, are structurally inadequate for governing neural data in such contexts — not because they are wrong, but because their individualist premise is misaligned with communities in which personhood is understood relationally.

Drawing on Ubuntu's relational ontology and the Zambian Chilimba communal pooling tradition, this article proposes the Neural Chilimba framework — a three-pillar governance model that operationalises four neuro-rights (cognitive liberty, mental privacy, mental integrity, and psychological continuity) at the community level. The framework operates within Gyekye's moderate communitarianism: it does not abolish individual parental consent but embeds it within communal deliberative governance, arguing that where data is inherently relational in its epistemic content, individual consent alone is a structurally inadequate governance instrument.

Legal analysis of Zambia's Data Protection Act No. 3 of 2021 identifies four structural gaps that leave neural data unprotected. Analysis of five consumer neurotechnology products, supported by a formal scoring rubric and inter-rater verification (reported in full in supplementary materials), reveals that zero products offer privacy policies in any Zambian indigenous language and none disclose compliance with Zambian law. Twelve phased and costed policy implications are provided for Zambia's Data Protection Commission, Ministry of Technology and Science, Ministry of Education, technology providers, and the African Union. The framework includes mandatory pre-deployment efficacy disclosure, a three-architecture data processing classification, longitudinal data retention limits, and youth advisory representation on governance committees. This article is an anticipatory governance intervention, not an empirical study of large-scale deployment in Zambia; it combines doctrinal legal analysis, a product audit, and normative institutional design to specify governance conditions that should be in place before educational neurotechnology becomes entrenched in Zambian school settings.

1. Introduction

1.1 The neuro-educational frontier

Portable electroencephalography (EEG) headsets — devices capable of generating correlates of attention, cognitive load, emotional arousal, and fatigue — are transitioning from clinical research settings into educational pilot programmes across North America, Europe, and East Asia [1,2]. These neuro-education tools promise a qualitative shift in personalised learning: real-time feedback on student engagement, [3]early identification of cognitive difficulties, and evidence-informed optimisation of pedagogical strategies [4]. Growth in the educational neurotechnology market has been substantial, with multiple peer-reviewed analyses projecting continued acceleration over the coming decade [5,6].

These developments are unfolding primarily in well-resourced educational systems with comparably mature data governance infrastructure. Their diffusion into less-regulated contexts — particularly the Global South — is underway but underexamined.[7] Consumer-grade EEG devices are available on Zambian e-commerce platforms; aid-funded digital-inclusion initiatives have piloted educational technology, including early-stage neurotechnology devices, in sub-Saharan African schools[8]; and corporate social responsibility programmes have offered donated consumer cognitive devices to schools in West and East African countries [1]. These schools frequently lack the legal expertise or regulatory protection to evaluate what accepting such devices entails. Privacy policies for these products are written exclusively in English; terms of service reserve the right to share 'de-identified' data with unspecified 'research partners'; and — as the product analysis in Section 4 documents — none of the five products reviewed disclosed compliance with Zambian national law.[7]

1.2 The problem: defining neuro-extractivism with precision

This article introduces the term neuro-extractivism to describe the asymmetric and largely irreversible flow of children's cognitive data from LMIC educational settings to Global North technology corporations. The term is an extension of Couldry and Mejias's [9] data colonialism framework. We demonstrate directly that it adds analytical value beyond existing concepts. Neuro-extractivism is not presented here as a categorically distinct phenomenon from data colonialism but as a regulatory intensification of it: three aggravating features combine to make educational neural data a governance problem that existing frameworks — however adequate for other data types — are not structurally designed to handle.

The first aggravating feature is irreversibility with compounding inference. All digital data, once incorporated into trained AI models, is practically non-deletable — this is not unique to neural data. What is distinctive is the compounding inference problem: consumer EEG devices do not read thoughts, and Muse-class headbands produce noisy, low-spatial-resolution signals with limited single-instance interpretability. However, when aggregated across hundreds of students and processed through commercial machine [10,11]-learning pipelines, EEG-derived attention and arousal proxies can generate population-level cognitive profiles that substantially exceed the interpretive content of any individual signal. Dwork et al.'s [12] foundational work on privacy amplification by aggregation demonstrates this general mechanism; Fairclough and Gilleade [13] specifically document the gap between raw biosignal data and the inferences enabled by its commercial processing. It is the aggregated inference, not the individual signal, that constitutes the primary governance risk, and it is against this inference layer that the Neural Chilimba's prohibitions and audit rights are directed.

The second aggravating feature is the involuntary disclosure of pre-reflective states — that is, cognitive states that occur without the subject's conscious awareness or intentional control, such as baseline arousal, attentional drift, or subconscious emotional responses to stimuli. Behavioural data extracted by learning management systems requires some form of active disclosure: a student chooses to type, click, or submit. EEG signals are generated involuntarily during normal cognitive activity. A student cannot choose not to produce neural signals while wearing a headband, any more than they can choose not to have a heart rate. Applying Nissenbaum's contextual integrity framework, EEG data collected in a school setting violates contextual norms because the norms of that context (supervised education, not biomedical research or commercial profiling) do not include involuntary cognitive data collection. This distinction from platform behavioural data — which Williamson [14] and Selwyn [15] have documented as exploitative — is not a difference of kind but of the degree of disclosure control available to the data subject, which is zero in the EEG case.

The third aggravating feature is the institutional vulnerability of the school. The school carries state sanction, parental delegation of care, and the developmental significance of cognitive formation. These are features that social media platforms, however dominant, do not possess. Neural data extracted from a school therefore occupies a qualitatively different position in the architecture of children's development, because the school is the site where cognitive patterns that EEG measures are being actively formed. This is not a claim that no other data governance framework has ever identified school data as sensitive — Lupton and Williamson [16] and Prinsloo and Slade [17] have documented schooling's role as a dataveillance site — but that the combination of involuntary disclosure and developmental site specificity creates an aggregation-and-inference risk that existing ed-tech governance literature, focused on behavioural data, does not yet fully address.

We do not claim neuro-extractivism requires an entirely separate theoretical edifice. We claim it requires regulatory provisions that existing data colonialism and surveillance capitalism frameworks have not yet generated: mandatory collective consent, pre-deployment efficacy disclosure, population-level inference prohibition, and benefit-sharing agreements tied to community governance structures. The Neural Chilimba provides these provisions.

1.3 The opportunity: indigenous governance as principled innovation

Rather than recommending that Zambia reject educational neurotechnology, a position that would be developmentally counterproductive and technologically unrealistic, this article argues that Zambia has resources, both philosophical and institutional, to build governance infrastructure that is more protective and more contextually grounded than frameworks produced in the Global North for different conditions.

The Ubuntu ethic; munthu ndi muntu pa anthu in Zambia's Nyanja language, 'a person is a person through people' offers a relational account of personhood that generates different governance conclusions than liberal individualism. This relational ontology mirrors the technical architecture of the biocybernetic loop and physiological computing ecosystems themselves, where cognitive data is not a static individual asset but an emergent property generated through continuous, recursive interaction between the user, the device, and the immediate environmental context [13,18]. The Chilimba tradition, a Zambian rotating savings and credit association (ROSCA) widespread in rural and peri-urban communities [19], provides a ready institutional logic for pooling communally generated resources under community governance. These are not abstract philosophical resources being applied to a novel problem; they are living institutional traditions that Zambian communities already use to manage collective stakes in shared resources.

The case for anticipatory governance is established by the history of digital technology regulation. Social media, facial recognition, and genomic data collection were all subject to governance frameworks developed after the extractive practices they were meant to constrain were already entrenched [20]. The Neural Chilimba is a precautionary framework in the precise sense: it builds governance infrastructure before extraction begins, not after.

1.4 Research approach, scope, and contribution

This article proceeds through three analytical phases. The first is legal document analysis of Zambia's Data Protection Act No. 3 of 2021 and related instruments, identifying the regulatory gaps that make educational neurotechnology ungovernable within the existing framework. The second is narrative literature synthesis across six bodies of scholarship, generating the conceptual foundations from which the framework is built. The third is the development and operationalisation of the Neural Chilimba framework itself, translated into twelve phased policy implications.

Three research questions guide the analysis: (1) What are the specific regulatory gaps in Zambia's Data Protection Act No. 3 of 2021 with respect to educational neurotechnology? (2) How can Ubuntu and Chilimba be operationalised into a governance framework for educational neural data that is philosophically coherent and institutionally realistic? (3) What phased and feasibility-tested policy implications emerge from this framework?

The contribution is twofold. This article provides one of the first context-specific analyses of neuro-educational technology governance risks in the Zambian legal and cultural setting. It proposes the Neural Chilimba as a candidate governance framework — Version 1 of a contextually grounded model that treats Zambian indigenous philosophy not as a normative supplement to Western frameworks but as a primary source of governance logic. The framework is offered as a provisional architecture requiring empirical validation and community co-design before policy translation, not as a finished model.

Several important scope limitations must be stated at the outset. EEG-based educational technology has not yet been deployed at scale in Zambian schools; the framework is therefore anticipatory in design. Consistent with the tradition of anticipatory ethics [21], we argue that the absence of deployment is the reason to act now, not a reason to wait [22]. The article does not include stakeholder consultations with Zambian communities or a full cost-feasibility analysis; both are priorities for the next phase of research [23]. This article is intentionally selective rather than prohibitive: it does not reject educational neurotechnology in principle, but conditions legitimacy on demonstrable efficacy, auditability, and rights-compatible processing architectures. To prevent any misreading of the paper's evidentiary scope, the following non-claims are stated explicitly: this article does not claim that educational neurotechnology is currently deployed at scale in Zambian schools; it does not present a field study, pilot implementation, or stakeholder co-design exercise; it does not claim that the proposed framework is implementation-ready in full under current Zambian state-capacity conditions; it does not provide a complete engineering blueprint for all processing architectures; and it does not claim that neuro-extractivism is wholly discontinuous from adjacent frameworks such as data colonialism, surveillance capitalism, or biometric governance. What this article does claim is that the specific configuration of involuntary pre-reflective disclosure, compounding population-level inference, and institutional vulnerability of the school as a site of cognitive formation constitutes a distinctive regulatory problem that existing frameworks have not yet specified governance provisions for.

2. Literature Review and Conceptual Background

2.1 Neuro-education: promises, evidence problems, and governance implications

The application of neuroscience to education has been a growing field since the early 2000s [24]. The advent of portable, low-cost EEG devices has accelerated this agenda, enabling classroom-based studies previously impossible in terms of scale and ecological validity [25]. The promises are real. Real-time cognitive-load monitoring can help teachers identify when students are overwhelmed and adjust instruction accordingly [26]. Attention-tracking can reveal which lesson components engage or lose students [18]. Emotional-valence detection can flag students experiencing distress for early intervention [16]. However, the evidential basis for these promises in real classroom contexts is considerably weaker than commercial accounts suggest, and this evidential weakness has direct governance implications. Howard-Jones [24] has documented the proliferation of neuromyths in educational settings — scientifically unfounded beliefs that persist because they draw on neuroscientific-sounding language. The replication crisis in psychology and cognitive neuroscience has demonstrated that many brain-behaviour relationships established in controlled laboratory settings do not translate reliably to the ecological complexity of real classrooms [27]. The relationship between EEG-derived metrics and meaningful learning outcomes remains contested even in well-resourced research settings [25]. Wexler [28] has specifically documented outcome-reporting bias in commercial neurotechnology claims, and Pasquinelli [29] provides the most comprehensive account of why neuromyths persist and spread. These evidential limitations require a dual-scenario governance approach. A governance framework for neuro-educational technology must be designed for two distinct cases: (A) the technology produces reliable correlates of cognitive engagement — in which case the data extraction risks are real and the governance structure must prevent exploitation of genuine signals; and (B) the technology produces unreliable, noisy, or commercially overstated signals — in which case the governance structure must prevent a different set of harms: false pedagogical labelling, misallocated resources, and stigmatisation of children based on meaningless metrics. The Neural Chilimba addresses both scenarios. Scenario A is addressed through the data sovereignty and benefit-sharing architecture; Scenario B is addressed through the mandatory pre-deployment efficacy disclosure requirement introduced in Pillar 3 (Section 5), which conditions community approval on independent peer-reviewed evidence of classroom validity in comparable contexts.

Almost all of this literature is produced from and for WEIRD (Western, Educated, Industrialized, Rich, Democratic) contexts [30], generating additional uncertainty about whether EEG-classroom findings would replicate in Zambian educational environments. The specific vulnerabilities of educational neurotechnology in the Global South — weaker regulatory enforcement, lower digital literacy, linguistic barriers to informed consent, and the absence of neuro-specific legal provisions — remain almost entirely unexplored in the peer-reviewed literature [31].

2.2 Data colonialism and neuro-extractivism as regulatory intensification

Couldry and Mejias [9] argue that contemporary data extraction constitutes a new phase of colonialism in which human life, rendered as data, is appropriated for capitalist accumulation without meaningful consent or compensation. Kwet [26] and Selwyn [15] and Williamson [14] document how ed-tech platforms extract vast quantities of learner data from schools in the Global South under the guise of development or digital inclusion. These analyses provide the foundational framework within which neuro-extractivism is situated.

Neuro-extractivism, as defined in Section 1.2, is an intensification of data colonialism along three dimensions: the irreversibility-plus-compounding-inference problem, the involuntary disclosure of pre-reflective states, and the developmental and institutional vulnerability of the school as extraction site. These are not categorical differences that require a wholly new theoretical framework; they are aggravating features that require regulatory provisions that the data colonialism literature has not yet specified. The Neural Chilimba provides those provisions.

The distinction from surveillance capitalism [20] is similarly one of degree and specificity. Zuboff's account addresses the extraction of behavioural surplus — clicks, searches, purchases — for prediction and influence. EEG-derived cognitive metrics differ in that they are generated involuntarily and encode cognitive states that precede any behavioural expression, including states the subject is not aware of [32,33]. The governance implications of pre-reflective data differ from those of behavioural data in one specific respect relevant to school contexts: they cannot be governed by the existing ed-tech consent and opt-out frameworks, because no meaningful opt-out from involuntary neural data generation is possible during the act of learning.

Table 1 (below) positions neuro-extractivism against five adjacent concepts to demonstrate why it constitutes a distinctive regulatory configuration rather than a mere relabelling of existing frameworks. The table does not claim categorical difference across all dimensions — only that the specific combination of features in column 6 has not been addressed as a unified governance problem by any existing framework.

Table 1 Why neuro-extractivism is not reducible to adjacent concepts. The table does not claim categorical difference across all dimensions — only that the specific combination of features in column 7 has not been addressed as a unified governance problem by any existing framework.

Concept	Primary extraction object	Mode of disclosure	Institutional site	Data subject vulnerability	Main governance assumption	Why insufficient for educational neural data
Data colonialism	Behavioural/transactional data	Active/passive digital trace	Platform	General	Consent and data rights	Does not specify collective consent or pre-deployment efficacy disclosure requirements
Surveillance capitalism	Behavioural surplus	Active user behaviour	Commercial platform	General	Prediction products prohibition	Does not address involuntary pre-reflective disclosure or community governance
Biometric governance	Physical identifiers	Voluntary enrolment or passive scan	Border/workplace/device	General	Special category consent	Does not address compounding inference from aggregate neural signals
Psychometrics	Trait-level inference	Voluntary testing	Educational/clinical	Assessed adults or consenting minors	Individual consent	Does not address real-time involuntary continuous collection in classroom settings
Genomics	Genetic sequences	Biological sample	Clinical/research	Research participants	Research ethics board + consent	Does not address classroom community relational context or community benefit-sharing
Neuro-extractivism (this article)	Involuntary pre-reflective EEG signals + population-level ML inference	Involuntary/continuous	School (state-sanctioned, developmental site)	Developing children	No adequate existing framework	Requires collective consent + pre-deployment efficacy disclosure + community benefit-sharing + processing architecture requirements

2.3 African data governance: regulatory landscape and indigenous resources

Forty-four African nations had enacted data protection laws by early 2026, most modelled closely on the EU's GDPR [34]. The African Union's Data Policy Framework (2022) [2]sought to provide a continental-level governance architecture, but like its national counterparts, it does not address neural data as a distinct category. Zambia's Data Protection Act No. 3 of 2021 [35] establishes a Data Protection Commission, requires consent for data processing, and grants data subjects rights of access and deletion. As demonstrated in Section 4, however, the Act has critical gaps with respect to neural data.

A growing body of scholarship has argued that indigenous African philosophies, particularly Ubuntu, offer resources for reimagining data ethics that Western frameworks structurally cannot provide [17,36]. The CARE Principles for Indigenous Data Governance [37]—Collective Benefit, Authority to Control, Responsibility, Ethics — represent the most developed international framework in this tradition. The OECD Recommendation on Responsible Innovation in Neurotechnology (2019) calls for adaptive governance of neuro-devices but does not specify the community-level institutional mechanism through which adaptation occurs. The Neural Chilimba operationalises the CARE Principles within the specific constraints of Zambian educational institutions and provides the institutional mechanism the OECD recommendation lacks.

2.4 Ubuntu as a foundation for relational data governance

Ubuntu, a normative philosophical tradition rooted in southern and central Africa and extensively analysed in the African philosophy literature [38], is expressed in Zambia's Nyanja language as munthu ndi muntu pa anthu — 'a person is a person through people'. While frequently invoked as a general ethic of care or communalism, Ubuntu denotes a specific ontological claim about personhood: that personhood is not prior to social relations but constituted through them. This ontological claim generates governance conclusions when applied to neural data collected in relational educational contexts. We develop this argument in three steps. The first step is ontological. Ubuntu holds that a person's cognitive dispositions, learning patterns, and identity are not pre-social but constitutively shaped by the relational matrix of family, teachers, peers, and community [29]. This is not an exotic claim; it is broadly consistent with developmental psychology's account of the social basis of cognitive development [39] and with the neuroscientific literature on brain-to-brain synchrony in educational settings [25,40].

The second step is epistemic. EEG data collected in a school context does not measure only individual cognitive states; it encodes relational cognitive patterns. Dikker et al.[25] and Bevilacqua et al.[40] demonstrate that classroom EEG signals reflect brain-to-brain synchrony between students and teachers, and that this synchrony correlates with learning outcomes. The data generated is therefore not purely about the individual student — it is inherently relational in its epistemic content, encoding the cognitive texture of a shared educational relationship.

The third step is normative (epistemic). Where data is inherently relational in its epistemic content, individual consent frameworks are structurally inadequate because they cannot account for the collective dimension of what is being measured [41]. An individual parent cannot consent on behalf of the relational cognitive environment from which the data derives. The community governance claim that follows is procedural rather than proprietary: it is not that the community owns the child's neural data, but that individual consent alone cannot adequately govern data whose epistemic content is irreducibly social. The fourth step is normative (disclosure). EEG data is generated involuntarily during normal cognitive activity: a student cannot choose not to produce neural signals any more than they can choose not to have a heart rate. Where disclosure is involuntary, the standard opt-in/opt-out consent architecture is insufficient as a complete governance instrument [42]. The absence of any disclosure control — which is zero in the EEG case — means that the governance burden cannot rest on individual consent alone; it must include a prior communal decision about whether the context of collection is appropriate at all. This is a condition that the community, not any individual parent acting alone, is best placed to evaluate and enforce. The fifth step is normative (vulnerability). Children at the stage of cognitive development represented by the school-age population are subjects of special protection under international law [43] and are incapable of providing meaningful consent to neural data collection on their own behalf. Where both the data subject (the child) and the data type (involuntary pre-reflective EEG) are simultaneously incapable of individual consent governance, communal oversight is not merely desirable but structurally necessary [43,44]. These three normative premises — epistemic relationality, involuntary disclosure, and child developmental vulnerability — jointly establish why individual consent alone is structurally inadequate: no single premise is sufficient alone, but their combination forecloses individual-consent governance as a complete solution. This procedural relational claim is consistent with, and does not require overriding, individual parental opt-out rights, a parent can withdraw their child from participation without negating the community's governance role over the remaining pool.

The procedural relational claim grounds the governance argument in the epistemic content of the data itself rather than in the metaphysics of personhood alone. The fact that a community influences cognitive development does not, without further argument, entail community governance authority over data measuring that development; the argument developed here avoids this inference gap by locating the governance claim in what the data measures, not in who shaped the person being measured. The present framework operates explicitly within Gyekye's [45] moderate communitarianism. Gyekye argues that communal frameworks must preserve individual rights, cautioning against Ubuntu interpretations that subordinate the individual to collective will. The Neural Chilimba honours this caution structurally: individual parental consent is a necessary condition for deployment, preserved within rather than displaced by communal governance. The community adds a governance layer; it does not remove individual protection.

2.4.1 Ubuntu's internal tensions: feminist and postcolonial critiques

Ubuntu cannot be treated as a normatively self-authorising framework. Oyèrónkẹ́ Oyěwùmí has argued that Ubuntu as canonised in African philosophy literature frequently encodes male-centred assumptions about communal authority, reflecting the androcentrism of the philosophical traditions that

produced its canonical formulations. Nkiru Nzegwu [46] demonstrates that African communal frameworks have historically been invoked in ways that domesticate women's agency and vest communal authority in male elders. These critiques directly inform the Neural Chilimba's governance design. The framework distinguishes between what we term deliberative Ubuntu — an interpretation that requires structured representation of multiple community voices, including youth, women, and linguistic minorities — and gerontocratic Ubuntu, which vests communal authority in elders by virtue of age or gender. The Neural Chilimba explicitly adopts the former and rejects the latter. Gender-balanced Committee composition (minimum 50% of elected seats held by women), youth advisory representation, and the prohibition on restricting the Community knowledge-keeper seat to male elders are non-negotiable structural requirements, not aspirational guidelines. The Data Stewardship Register's public transparency mechanism ensures that any drift toward gerontocratic governance is visible and contestable.

2.5 Chilimba as data pooling logic: three-level analysis

Chilimba (Chilimba chathu — 'Our resilience') is a Zambian rotating savings and credit association (ROSCA) prevalent in rural and peri-urban communities [32]. Members contribute money, labour, or agricultural produce to a common pool, which is then distributed to members in need at zero interest. The system operates on three principles: collective contribution, rotational benefit, and trust-based accountability. Our proposal adapts the institutional logic of Chilimba to the domain of educational neural data governance. To be clear about where the analogy holds and where it requires supplementation, we distinguish three levels:

Level 1: Institutional logic (holds). Collective contribution to a shared pool, governed by community members rather than external parties, with agreed-upon rules for benefit distribution and accountability mechanisms for non-compliance. This level of the analogy is structurally sound. It does not depend on the fungibility of the pooled resource; it depends on the principle of communal governance of a collectively generated asset. Neural data, collectively generated through communal educational activity, qualifies as such an asset [41]. This is the level at which the Chilimba logic makes its most important contribution: it provides an indigenous institutional template for the Data Chilimba Committee and its governance authority.

Level 2: Rotation mechanism (partially holds). In a traditional Chilimba, the pool itself rotates — each member takes a turn receiving the collected fund. Neural data does not rotate in this sense, and we do not claim it does. What rotates is the prioritisation of benefit allocation: the Committee's authority to determine which cohort's data insights are acted upon first, which pedagogical dividend is realised in which term, and which aspect of school infrastructure receives the annual contractual dividend from the provider. The equity logic — no cohort is permanently a net contributor without receiving corresponding benefit — is preserved through rotation of benefit priority, not rotation of data.

Level 3: Trust enforcement (does not hold; supplemented by contract). Traditional Chilimba governance operates through interpersonal, reputationally enforced trust among known community members [32,28]. A technology provider is not a Chilimba member. The contractual supplements that bridge this gap — benefit-sharing agreements, right-to-audit clauses, Committee oversight powers, anti-extraction prohibitions — are regulatory instruments, not relational governance mechanisms [41]. The Neural Chilimba is therefore a hybrid: relational governance within the school community, contractual governance at the school-provider interface. We state this explicitly rather than concealing it: the Chilimba logic applies to the community's internal governance, not to the community-corporation relationship, which is governed by contract and law.

2.6 Mapping the Neural Chilimba to established neuro-rights

Ienca and Andorno [47] and Yuste et al. [42] propose four foundational neuro-rights: cognitive liberty (freedom from mental manipulation), mental privacy (protection of neural data from unauthorised access), mental integrity (protection against unauthorised alteration of cognitive states), and psychological continuity (protection of cognitive identity over time). The legal grounding of cognitive liberty has been developed by Bublitz [10] and Farahany [44], whose account of mental self-determination directly informs the Neural Chilimba's prohibition on manipulation AI use.

Cognitive liberty maps to Mandate 3.3, which prohibits the use of Zambian children's neural data in commercial or political manipulation AI systems. Mental privacy maps to Mandate 3.1 (processing architecture requirements) and to the collective consent mechanism in Pillar 1. Mental integrity is addressed through the benefit-review cycle in Pillar 2, through which the Committee may revoke consent if a deployment influences student cognitive development in unauthorised ways. Psychological continuity — the protection of cognitive identity from longitudinal profiling — is addressed by Mandate 3.5, which prohibits retention of identifiable neural data beyond the current academic year and prohibits cross-cohort longitudinal linking that would enable cognitive identity profiling. This mandate is operationalised in Pillar 3 and summarised in Table 3a.

3. Analytical Methodology

3.1 Research design: anticipatory framework development

This article employs a two-stream analytical methodology appropriate for conceptual framework development in emerging technology governance contexts. The first stream is legal document analysis of Zambia's Data Protection Act No. 3 of 2021 and related instruments. The second is narrative literature synthesis across six bodies of scholarship. Together, these streams generate the empirical and conceptual foundations from which the Neural Chilimba framework is developed.

The design follows the tradition of anticipatory ethics [14]: rather than waiting for harm to occur before constructing governance, it builds normative infrastructure in advance of deployment — a choice strongly supported by the history of poorly governed technology adoption. The framework is anticipatory in the sense that it targets a deployment scenario that has not yet occurred at scale in Zambia. This requires a precise distinction: the article does not claim that neuro-extractivism is currently occurring in Zambian schools; it claims that the conditions that would make it possible if technology deployment occurred are present and documented. The distinction between risk documentation and harm documentation is maintained throughout.

The dual-scenario design principle follows from this. Anticipatory governance must account for the full range of credible deployment outcomes, including scenarios in which the technology does not work as commercially claimed (see Section 2.1). The pre-deployment efficacy disclosure requirement introduced in Pillar 3 is the governance mechanism that addresses Scenario B (technology fails to produce valid signals), ensuring that community approval is conditional on demonstrated validity rather than commercial assertion.

3.2 Stream 1: Legal document analysis

Five categories of primary source material were analysed. First, Zambia's Data Protection Act No. 3 of 2021 was analysed in full, with attention to Part IV (sensitive personal data), Section 14 (special categories), Section 17 (child data provisions), and the enforcement and penalty provisions. Second, the Zambia Demographic and Health Survey 2021–2022 [48], the 8th National Development Plan (2022–2026) [49], and the National Digital Transformation Strategy (2023–2027) [50] were examined for contextual and fiscal analysis. Third, five consumer neurotechnology products available on Zambian e-commerce platforms as of December 2025 were selected and analysed (detailed in Section 4.2 and Supplementary Materials). Fourth, four international governance frameworks were examined comparatively. Fifth, the Data Protection Commission's publicly available annual reports and press releases were reviewed to assess enforcement capacity.

A threshold legal question is whether neural data might already be covered as biometric data under Section 14 of the Act. Three considerations argue against relying on that interpretation. First, the DPC has not issued guidance confirming this classification. Second, consumer neurotechnology companies have successfully resisted biometric classification in EU, US, and Australian regulatory contexts by invoking the 'wellness device' distinction [28,20]; the same argument is available in Zambia. Third, any gap between deployment and judicial or regulatory clarification would leave children's neural data in an unprotected interim period. An explicit neural data addendum eliminates this ambiguity. This question is addressed in detail in Section 4.1.

3.3 Stream 2: Narrative literature synthesis

The narrative synthesis drew on six bodies of scholarship: (1) educational neuroscience and neuro-education governance [1,3,40,14,24,26,16,24,29,51]; (2) data colonialism and surveillance capitalism [9,20,26]; (3) African data governance and Ubuntu philosophy [7,1,1,13,18,2,17,36,41,46,7,44]; (4) indigenous data governance [7,31,37]; (5) neuroethics and neuro-rights [4,28,44,46,20,10]; and (6) anticipatory ethics and technology governance [13,21-23,47].

Searches were conducted in PubMed, PhilPapers, Google Scholar, and the African Journals Online (AJOL) database, supplemented by a targeted French-language search in AJOL covering Francophone African data governance scholarship. Search terms included: ('neuroeducation' OR 'educational neuroscience') AND ('governance' OR 'regulation'); ('EEG' OR 'neurotechnology') AND ('Global South' OR 'Africa' OR 'Zambia'); ('Ubuntu' OR 'African philosophy') AND ('data ethics' OR 'privacy'); ('data colonialism' OR 'neuro-rights' OR 'cognitive liberty'). Searches were limited to publications from 2010 to 2025, with exceptions for foundational texts. A supplementary French-language search in AJOL used the terms (‘gouvernance des données’ OR ‘éthique numérique’) AND (‘Afrique’ OR ‘Ubuntu’), conducted in November 2025; this search identified three Francophone African data governance sources that have been incorporated where relevant. A limitation of the English-and-French focus remains: Ubuntu and Chilimba scholarship in Nyanja, Bemba, and other indigenous languages was not retrieved by either search. Future research should address this through community-based and ethnographic methods.

4. Legal Document Analysis: Regulatory Gaps and Product Findings

4.1 Zambia's Data Protection Act No. 3 of 2021: four structural gaps

The legal document analysis identified four structural gaps in Zambia's Data Protection Act No. 3 of 2021 with respect to educational neurotechnology. These gaps are not incidental drafting oversights; they reflect the foundational design assumption of the Act — shared with GDPR-derived legislation generally — that data protection is primarily an individual rights matter, not a communal governance question.

Table 1a Structural gaps in Zambia's Data Protection Act No. 3 of 2021 and proposed Neural Chilimba responses.

Gap	Description — Data Protection Act No. 3 of 2021	Neural Chilimba Response
Gap 1: Neural data absent from sensitive-data categories	Part IV and Section 14 enumerate health data, biometric data, and genetic data as special categories. Neural data is not mentioned. Consumer EEG devices are marketed as wellness products, enabling providers to self-classify outside the heightened protection regime.	Neural Data Addendum (Implication 1) classifies all EEG-derived cognitive data as a sensitive special category under Section 14 regardless of device medical designation. Burden of proof shifts to the provider.
Gap 2: Consent assumes literacy and language access	Product analysis found all five privacy policies available exclusively in English at post-secondary reading level. Primary-education completion rates fall below 40% in rural Zambia [48]. A policy that cannot be read cannot produce genuine informed consent.	Mandate 3.2 requires all consent materials in English and at least two local languages (Bemba, Nyanja, Tonga, or Lozi) at Grade 7 reading level. A Community Glossary of neuroscientific terms, co-authored with community members, is mandatory.
Gap 3: Enforcement capacity structurally inadequate	The DPC operates with limited staffing and no demonstrated neuro-forensic expertise. Penalties are insufficient to deter multinational corporations for whom fines represent a cost of business rather than a deterrent.	Phased capacity-building: Year 1–3 secondment of DPC staff through University of Zambia Law School cyber-law clinic; Year 3–5 establishment of three-person neuro-forensics unit. Right-to-audit clause mandatory in all procurement contracts.
Gap 4: No provision against cognitive manipulation	The Act does not address the use of neural data to train persuasion or micro-targeting AI. Data collected in Zambian classrooms could be used to train political or commercial manipulation systems operating elsewhere.	Mandate 3.3 explicitly prohibits use of Zambian educational neural data in political or commercial manipulation AI. All secondary research use requires pre-approval by the Data Chilimba Committee and the DPC, with binding benefit-sharing agreements.

The question of whether neural data might already be covered as biometric data under Section 14 warrants direct engagement. Three considerations argue against relying on existing biometric classification alone. First, the DPC has issued no guidance confirming this classification. Second, consumer neurotechnology companies have successfully resisted biometric classification in EU, US, and Australian regulatory contexts by invoking the 'wellness device' distinction [28,20]; the same argument is available in Zambia and would need to be resolved by litigation or regulatory guidance before providing reliable protection. Third, any gap between deployment and judicial or regulatory clarification would leave children's neural data unprotected during the interim period. An explicit Section 14 addendum eliminates this ambiguity.

4.2 Consumer neurotechnology product analysis

Five consumer neurotechnology products available on Zambian e-commerce platforms as of December 2025 were selected using three criteria: marketed for attention, focus, or cognitive improvement; purchasable without a medical prescription; and capable of EEG or equivalent neural data collection. Full product names, privacy policies, terms of service, data-sharing provisions, and retrieval dates are documented in Supplementary Data File 1 (SDFile_1). Products are identified in the main text by type-class rather than brand name given ongoing legal proceedings in other jurisdictions against terms-of-service criticism in academic publications; the supplementary file enables complete replication of the analysis.

The Sovereignty Risk rating (Critical/Moderate) in Table 2 is assigned using a formal three-dimension rubric documented in Supplementary Methods File 1 (SDFile_2): (1) data transmission architecture (on-device processing to mandatory cloud, rated 1–3); (2) third-party data sharing scope (none stated to broad and unspecified, rated 1–3); and (3) local storage availability (default local to no local option, rated 1–3). Critical = composite score 7–9; Moderate = 4–6. A second rater (a legal colleague with data protection expertise) independently applied the rubric; inter-rater agreement was 4/5 products (80%). The single disagreement (Product C) is reported in SDFile_2.

Table 2. Consumer neurotechnology products available in Zambian markets (December 2025). *Inter-rater disagreement on Product C documented in SDFile_2. Full product documentation in SDFile_1.

Product	Type	Data Collected	Policy Languages	Third-Party Sharing	Local Storage	Sovereignty Risk
A (Muse-class EEG headband)	Attention/focus	Raw EEG; processed attention metrics	English only	Yes ('aggregated')	No	Critical (8/9)
B (Neurofeedback app + headset)	Affective computing	Processed emotional valence; fatigue indices	English only	Yes ('de-identified')	No	Critical (7/9)
C (Meditation headband)	Mindfulness EEG	Raw EEG; meditation depth score	English, French	Yes ('research partners')	No	Critical (7/9)*
D (Cognitive training headset)	Cognitive load	Processed cognitive load index	English only	No (stated)	Partial (optional cloud)	Moderate (4/9)
E (Student focus wearable)	Classroom attention	Attention level; distraction duration	English only	Yes (unspecified)	No	Critical (8/9)

Several findings from this analysis warrant emphasis. First, zero of the five products offered privacy policies in any Zambian indigenous language. Second, four of five products reserved the right to share neural data with third parties using language — 'aggregated,' 'de-identified,' 'research partners' — that is deliberately vague about the nature and beneficiaries of sharing. Third, only one product offered any local storage option, and this was optional rather than default. Fourth, none of the products' terms of service referenced the Zambian Data Protection Act No. 3 of 2021. These findings concern the five products reviewed and are not presented as representative of the global neurotechnology market; they are consistent with, and provide Zambia-specific documentation of, the broader patterns documented in the student data privacy literature.

The product analysis also reveals a specific legal vulnerability: Product D's self-declaration that it does not share with third parties cannot be independently verified, since there is no audit mechanism in the current DPC framework to test such declarations against actual data flows. The right-to-audit clause proposed in Mandate 3.4 is the governance response to this verification gap.

Note on the biometric classification question: No product reviewed disclosed that its EEG data was subject to the Zambian Act's biometric data provisions, and no product's documentation acknowledged that neural data might qualify as a special category. This is consistent with the industry-wide wellness-device classification strategy documented by Farahany [44] and Bublitz [10] in other jurisdictions.

5. The Neural Chilimba Framework

Table 3. Authority and Decision Matrix for the Data Chilimba Committee. Rows specify governance actions; columns specify decision roles.

Governance action	Parent / guardian	Child / youth representative	Data Chilimba Committee	School administration	Ministry / DPC / Clearinghouse
School-wide deployment eligibility	Vote (PTA supermajority)	Advisory — recorded in minutes	Authorise (75% supermajority)	Consult	Authorise (DPC registration)
Individual child participation	Authorise 2014; may withdraw at any time without penalty	Consult	Consult	Inform	2014;
Secondary data use approval	Inform	Advisory — recorded in minutes	Authorise 2014; renewed review required	Inform	Authorise (DPC pre-approval)
Contract renewal / renegotiation	Inform	Advisory — recorded in minutes	Authorise	Consult	Inform
Dispute: inferential harm to opted-out child	Complainant	Advisory	First instance review	Consult	Final determination (DPC)
Community complaint against provider	Complainant	Advisory	Investigation and recommendation	Consult	Enforcement (DPC / Ministry)
Emergency suspension of deployment	May request	May request	May recommend	May implement pending review	Authorise 2014; suspension authority
Audit escalation to DPC	May trigger	May trigger	Initiates escalation	Consult	Conducts audit; issues binding directions

Key: Authorise = decision-making authority; Consult = input required before decision; Inform = must be notified; Advisory = participates, views formally recorded but no veto. (1) PTA supermajority authorises school-level deployment eligibility but cannot compel individual participation. (2) Parental consent required and may be withdrawn at any time without penalty. (3) Neither PTA nor school may authorise secondary use beyond approved communal purpose; renewed Committee review required. (4) DPC and Ministry retain suspension authority where audit findings or rights violations warrant it. (5) Youth representatives hold advisory status; objections must be formally recorded in Committee minutes.

Table 3a. Neural Chilimba framework: three-pillar summary.

Pillar	Philosophical Grounding	Operational Mechanism	Neuro-Right Addressed
Pillar 1: Ubuntu as Relational Data Ethic	Ubuntu's procedural relational claim: where data is inherently relational in its epistemic content, individual consent alone is structurally inadequate. Moderate communitarianism (Gyekye [45]) preserves individual rights within communal governance.	Data Chilimba Committee holds collective governance authority. Individual parental consent is necessary but not sufficient; ≥75% PTA supermajority required. Individual opt-out preserved. Youth advisory seats (age ≥13) required. Gender balance non-negotiable.	Mental privacy [47]; Cognitive liberty [10]
Pillar 2: Chilimba as Data Pooling Mechanism	Chilimba ROSCA logic at Level 1 (institutional governance) and Level 2 (rotational benefit equity). Level 3 (trust enforcement) supplemented by contractual mechanisms, not claimed as analogous.	Neural data pooled as aggregated, anonymised metrics. Benefit allocation prioritised by cohort term. Annual contractual dividend from provider. Committee reviews and renegotiates. Data Stewardship Register publicly posted within 14 days of each meeting.	Mental integrity [47] via benefit-review cycle
Pillar 3: Neuro-Responsible Governance	Anticipatory ethics [11]; data sovereignty; dual-scenario design accounting for both technology efficacy and technology failure harms.	Three-architecture processing classification (see 5.3). Linguistic accessibility mandated. Pre-deployment efficacy disclosure required. Prohibition on manipulation AI use. Longitudinal data retention limited to one academic year. Youth voice in Committee.	Cognitive liberty; Psychological continuity [Mandate 3.5]

5.1 Pillar 1: Ubuntu as relational data ethic

Existing neuroethics frameworks treat neural data as individual property to be protected by individual consent. In Ubuntu contexts this is philosophically insufficient not because individual consent has no value — it does, and the Neural Chilimba preserves it — but because it is structurally incomplete as a governance instrument for data that is inherently relational in its epistemic content.

The Data Chilimba Committee holds collective governance authority over neural data deployed in its school community. Its composition — elected parents (50% minimum women), elected teachers (50% minimum women), one community knowledge-keeper (not restricted by gender or age), two youth advisory representatives (elected by peer process, gender-balanced, age ≥13 where applicable), and one non-voting technical representative — is designed to represent the community's diverse interests while insulating against domination by any single group.

Consent is layered. Individual parental consent is necessary but not sufficient for deployment; a ≥75% PTA supermajority approval is also required. Neither alone is sufficient; both together are necessary. Individual opt-out rights are preserved and are implemented silently: where a parent opts their child out, the child's device does not transmit data, and the school may not disclose which families have opted out or distinguish opted-out children during sessions through seating, equipment, or any other visible mechanism. This silent opt-out provision addresses the social-pressure problem inherent in communal approval contexts: formal opt-out rights are meaningless if exercising them invites social coercion. The framework acknowledges one irreducible visibility limitation: headsets are physically visible in the classroom, and peers will observe which children are wearing them. Silent opt-out therefore applies to formal institutional disclosure and data transmission, not to the physical visibility of devices. The framework prohibits institutional practices that would amplify or exploit peer awareness of device-wearing — such as seating arrangements, grade-book marking, or public acknowledgement — but cannot eliminate peer knowledge arising from ordinary classroom observation. Silent opt-out is therefore an administrative and anti-penalisation mechanism, not a guarantee of phenomenological invisibility: it minimises official exposure and coercive differentiation, but does not eliminate the social reality that peers can see who is and is not wearing a device. The framework's strength lies not in claiming to eliminate this residual visibility, but in making deployment conditional where social or inferential harms cannot be acceptably mitigated.

The aggregation inference problem also requires explicit acknowledgement. Even where a child's device does not transmit data, that child's cognitive environment — their interactions with teachers and peers — contributes to the relational signals that other students' devices may capture. This is an irreducible limitation of classroom EEG governance that no consent architecture can fully resolve. The Neural Chilimba treats it as an argument for the pre-deployment efficacy disclosure requirement: if a provider cannot demonstrate that opted-out children's cognitive signatures are not inferable from aggregate classroom data, that constitutes a ground for the Committee to deny deployment approval.

The ethical criterion for evaluating data use is communal flourishing, not merely individual privacy compliance. Disputes about data use are resolved through the Committee, which holds collective interpretive authority over whether a proposed use aligns with the community's endorsed purposes. Decisions are recorded in the Data Stewardship Register, posted publicly within fourteen days, and submitted annually to the DPC. 5.1.1 The child as rights-bearing subject within relational governance Relational governance in this framework is a supplement to, not a substitute for, the child's status as a rights-bearing subject. The Neural Chilimba does not reduce the child to a communal signal node or treat children's neural data as communal property. Four child-centred protections are non-waivable by any communal body, however constituted: (1) No communal body may authorise the use of a child's neural data for manipulation, behavioural micro-targeting, or commercial persuasion AI — the prohibition in Mandate 3.3 applies regardless of community approval. (2) No communal body may authorise indefinite retention of a child's identifiable neural data — the one-year deletion requirement in Mandate 3.5 is non-negotiable. (3) No communal body may authorise secondary commercial use of neural data beyond the original approved purpose without a new, distinct governance approval process. (4) Where a provider cannot demonstrate that opted-out children are protected against

inferential reconstruction from aggregate classroom data, deployment may not proceed — even with unanimous community approval. These provisions exist because communal governance can itself be a site of coercion, pressure, or miscalibrated interest. The child's rights under international law, particularly the Convention on the Rights of the Child's Articles 3 (best interests), 12 (right to be heard), and 16 (privacy) — set the floor below which no communal decision may reach. Youth advisory representation on the Committee ensures children's perspectives are structurally present; the non-waivable protections ensure they are structurally enforceable.

5.2 Pillar 2: Chilimba as data pooling mechanism

The Chilimba pooling logic fills the institutional gap that rights declarations alone leave open. Existing neuro-rights frameworks — Chile's constitutional amendment, the emerging neuro-rights literature — provide rights statements without specifying who enforces those rights at the community level, or through what mechanism benefit flows back to the communities whose data enables commercial value creation.

The Data Chilimba Committee governs the pooled data under annual rotation of elected seats to prevent entrenchment and co-optation. The Data Stewardship Register provides transparency: all Committee decisions, benefit-sharing negotiations, data-sharing approvals, and renewal outcomes are publicly posted. To further guard against elite capture — a documented risk in communal governance generally — the Register is also submitted annually to the DPC, creating an external accountability mechanism.

The rotation of benefit allocation operates on a cohort-term basis. Grade 4's attention data in Term 1 informs scheduling adjustments for the whole school. Grade 5's cognitive-load data in Term 2 informs teacher professional development. Grade 6's engagement data in Term 3 informs resource allocation. The technology provider supplies an annual contractual dividend — device replacements, consumables, or a funded teaching position — as part of the benefit-sharing agreement negotiated before deployment. It is a condition of any benefit-sharing arrangement that the dividend may not function as an inducement that undermines the voluntariness of community consent or creates financial pressure on schools or families to accept deployment: benefit-sharing is a consequence of legitimate governance, not a mechanism for securing it. The Committee reviews this dividend annually and may renegotiate if commercial value extracted by the provider is disproportionate to community benefit received. To be precise about the analogy: this is benefit rotation in the equity sense — no cohort is permanently a net contributor — not data rotation in the financial sense.

At the provider-relationship level, the Neural Chilimba operates through contractual law, not communal trust. Benefit-sharing agreements, right-to-audit clauses, and anti-extraction prohibitions are regulatory and contractual instruments. The Chilimba analogy applies to the school community's internal governance; the school-provider interface is governed by contract.

5.3 Pillar 3: Neuro-responsible governance

Pillar 3 provides the technical and legal mandates that translate the framework's philosophical commitments into enforceable governance. Its central technical mechanism is a three-architecture processing classification that addresses a fundamental property of consumer EEG systems: raw EEG data is meaningless noise without the provider's proprietary processing pipelines, and mandating local storage does not halt extraction if processing still requires transmission to proprietary cloud infrastructure.

The three processing architectures, with corresponding governance requirements, are:

Architecture A (On-device processing): All signal processing occurs on the headset or a local school device. No raw EEG leaves the physical school environment. Permitted as the default architecture upon standard Committee approval. Provider market preference is directed toward this architecture through procurement requirements: only Architecture A products are eligible for Ministry-level procurement contracts without additional conditions. This creates a market incentive for providers to develop on-device processing capabilities if they want access to Zambian educational markets.

Architecture B (Sovereign edge processing): Raw EEG is processed on a Ministry-managed edge server using open-source signal processing pipelines (MNE-Python, EEGLAB, or equivalent). Only derived metrics — not raw EEG — are transmitted for provider validation. Permitted with enhanced Committee approval and mandatory annual technical audit. Infrastructure managed by the Ministry of Technology and Science through a Service Level Agreement with the University of Zambia Computer Science Department.

Architecture C (Cloud-dependent processing): Raw or minimally processed EEG is transmitted to provider cloud infrastructure. Conditionally prohibited: permitted only where the provider grants a binding Zambia-processing clause requiring that all computation occurs on servers physically located in Zambia or on AU-compliant regional infrastructure, with right-to-audit of the processing environment. Where this cannot be provided, the product fails National Neuro-Ethics Clearinghouse certification. This effectively excludes most current consumer EEG products unless providers invest in local processing. That exclusion is a design feature: it creates a market incentive for Architecture A development.

Mandate 3.5 (Psychological Continuity Protections — new): Two provisions address the risk of longitudinal cognitive profiling. First, all identifiable or re-identifiable neural data must be securely deleted at the end of each academic year; deletion must be certified by the Committee and confirmed to the DPC in the annual governance report. Second, anonymised aggregate metrics may not be linked across student cohorts in ways that would enable longitudinal cognitive profiling of specific individuals or demographic subgroups. These two provisions together prevent construction of multi-year cognitive identity dossiers while preserving the pedagogical value of aggregate, cross-cohort insights at the population level.

Mandate 3.6 (Pre-deployment efficacy disclosure — new): Before the Committee may approve deployment, the technology provider must submit independent peer-reviewed evidence of classroom validity of its cognitive metrics in contexts comparable to the intended Zambian deployment. Evidence from WEIRD contexts [30] that has not been replicated in LMIC or sub-Saharan African school settings is insufficient. This condition implements the dual-scenario governance principle: approval requires demonstrated efficacy, not commercial assertion. Where no valid evidence exists, deployment is deferred pending independent trial. It is anticipated that, given the current state of the literature, most consumer EEG products will be unable to meet the Mandate 3.6 threshold in the near term; this is the framework's intended gatekeeper function. The governance posture during this interim period is one of precautionary deferral: no deployment proceeds without valid evidence. During the interim, the Phase 1 legislative and regulatory actions (Implications 1, 2, 6, and 8) remain fully operative as preparatory measures, and the National Neuro-Ethics Clearinghouse is established to receive and evaluate efficacy submissions as they become available. For completeness, the framework is selective rather than absolutely exclusionary: a product would be conditionally admissible if it met the following minimum combination of properties. It would need to demonstrate Architecture A (on-device) or auditable Architecture B (sovereign edge) processing; supply independent peer-reviewed evidence of classroom-validity in a sub-Saharan African or comparable LMIC school context; provide all user-facing materials in at least two Zambian indigenous languages at Grade 7 reading level; disclose the full scope of any third-party data sharing; grant the Data Chilimba Committee a contractual right-to-audit of data flows; and offer an annual contractual dividend as a condition of benefit-sharing. No current consumer EEG product known to the authors meets all six conditions, which accounts for the framework's near-term gatekeeper effect. But the conditions are not designed to be permanently unattainable: they are designed to be met by providers who invest in on-device processing, LMIC-context research, and genuine community benefit-sharing. That is the distinction between a selective governance framework and a disguised prohibition.

5.4 Comparative positioning

Table 4. Comparative positioning of the Neural Chilimba framework. A 'Known Weaknesses' column is included to provide an honest rather than self-serving comparison.

Framework	Jurisdiction	Legal Status	Consent Model	Distinction from Neural Chilimba	Neural Chilimba's Known Weaknesses vs. this Framework
Chile Neuro-Rights (Art. 19)	Chile	Constitutional (2021)	Individual rights-based; state protection	Individualistic; no communal governance body, benefit-sharing, or data localisation.	Chile's rights have binding constitutional force; Neural Chilimba is a proposed framework. Chile's implementation experience is an important learning resource.
EU AI Act (2024)	EU	Binding regulation	Risk-based classification; no neural-specific consent	Sector-wide AI governance; no community-level institution; not designed for LMIC contexts.	EU AI Act has mature enforcement infrastructure and broader jurisdictional reach. Neural Chilimba cannot replicate this scale.
UNESCO AI Ethics Recommendation (2021)	Global	Non-binding soft law	Human-rights framing; national implementation	Non-binding; no school-level institutional mechanism; no enforcement.	UNESCO framework has broader international legitimacy as an established standard; Neural Chilimba is untested and Zambia-specific.
CARE Principles for Indigenous Data Governance	Global (indigenous)	Ethical standard / soft law	Collective Benefit, Authority to Control, Responsibility, Ethics	Designed for research contexts, not educational technology deployment. Lacks processing architecture mandates, benefit-sharing rotation, or enforcement through national law.	CARE Principles have substantially broader international legitimacy and uptake. Neural Chilimba is an operationalisation for one specific context, not a rival standard.
Neural Chilimba (this article)	Zambia (proposed)	Proposed policy framework	Collective (Ubuntu-informed, procedural relational) + individual	Integrates communal governance, three-architecture processing classification, benefit-sharing rotation, efficacy disclosure requirement, longitudinal retention limits, and youth voice.	Untested; not piloted; not yet legally enacted; no stakeholder consultations completed; no cost-feasibility analysis; enforcement depends on DPC capacity that does not yet exist.

6. Normative and Institutional Implications

The following twelve implications derive from the legal analysis, literature synthesis, and Neural Chilimba framework. They are aligned with Zambia's 8th National Development Plan (2022–2026) and the National Digital Transformation Strategy (2023–2027). They are organised into three implementation phases based on resource requirements and dependencies: Phase 1 (Year 0–1, requiring no new capital investment); Phase 2 (Year 1–3, requiring capacity-building investment); and Phase 3 (Year 3–5, requiring capital and infrastructure investment). Scoping cost estimates, drawn from GIZ Digital Transformation in Africa programme benchmarks, are provided where available.

6.1 Phase 1 (Year 0–1): Legislative, regulatory, and policy actions

Implication 1 (Legislative amendment — Data Protection Commission): Propose an amendment to the Data Protection Act No. 3 of 2021 adding a Neural Data Addendum with four clauses: (a) neural data is classified as a sensitive special category under Section 14, regardless of device medical designation, with burden of proof on the provider; (b) collective as well as individual consent is required for educational neural data collection from minors; (c) neuro-extractivism as defined in this article is a prohibited data processing practice; and (d) benefit-sharing agreements are mandatory conditions of approval for any secondary use of educational neural data. The Addendum shall be reviewed every three to five years. Estimated cost: nil. Timeline: DPC to propose amendment to the Ministry of Justice within six months of framework adoption.

Implication 2 (Regulatory guidance — Data Protection Commission): Issue binding regulatory guidance applying the Act to consumer neurotechnology in educational settings, incorporating the three-pillar Neural Chilimba principles. Guidance must be translated into Bemba, Nyanja, Tonga, and Lozi at Grade 7 reading level and disseminated to all registered schools. Guidance must include gender-sensitive language and examples, and require intergenerational consultation in governance and consent processes. Estimated cost: USD 15,000–25,000 for translation and dissemination.

Implication 6 (Language mandate — Ministry of Technology and Science): Issue a procurement directive requiring that all educational technology acquired by or donated to government schools provide privacy policies, consent materials, and documentation in English and at least two local languages at Grade 7 reading level. Non-compliance voids procurement approval. Donated devices must be refused or quarantined pending compliance. Estimated cost: nil.

Implication 8 (School governance guidelines — Ministry of Education): Issue national guidelines for establishing Data Chilimba Committees at schools piloting neuro-educational technology. Guidelines must include: template terms of reference; compositional requirements (minimum 50% women in elected seats, youth advisory seats for students aged ≥13, community knowledge-keeper unrestricted by gender or age); annual rotation of elected seats; model benefit-sharing agreements; Data Stewardship Register requirements (decisions posted publicly within 14 days, submitted annually to DPC); silent opt-out procedures; and formal conflict-resolution pathways. Estimated cost: nil.

6.2 Phase 2 (Year 1–3): Capacity-building

Implication 3 (Capacity building — Data Protection Commission): Phase 2a (Year 1–2): Embed two DPC staff in the University of Zambia Law School cyber-law clinic as secondees, providing training in data protection enforcement, neuro-specific data classification, and digital forensics. Estimated cost: USD 20,000–30,000. Phase 2b (Year 3): Establish a minimum three-person neuro-forensics unit within the DPC in partnership with the University of Zambia and Copperbelt University, funded through the 8NDP Digital Transformation budget, with gender balance and succession planning requirements. Estimated cost: USD 80,000–120,000 per year at full operation.

Implication 7 (Teacher training — Ministry of Education): Develop a continuing professional development module on neuro-responsible technology use covering: EEG and neural data basics; Neural Chilimba principles; privacy policy assessment; and aggregated cognitive insight use without accessing individual student data. Piloted in urban flagship schools in Year 2; rolled out nationally through existing CPD infrastructure in Year 3. Estimated cost: USD 30,000–50,000 for development; marginal thereafter.

Implication 9 (Privacy by design — technology providers): Architecture A (on-device processing) is the required default and the only architecture eligible for government procurement without additional conditions. Architecture C (cloud-dependent) is conditionally prohibited absent a binding Zambia-processing clause. All user-facing materials must be available in Bemba, Nyanja, Tonga, and Lozi at Grade 7 reading level. Providers that cannot meet these requirements are excluded from the National Neuro-Ethics Clearinghouse registry and ineligible for government procurement.

Implication 10 (Benefit-sharing as market standard — technology providers): Compliance with the Neural Chilimba framework — genuine benefit-sharing, community governance participation, and anti-extractive data practices — is a condition for market access. Providers must: supply an annual contractual dividend before deployment approval; undergo annual independent audits of benefit-sharing effectiveness; and provide a formal community grievance mechanism with a maximum 30-day response requirement. Non-compliant providers are excluded from the Clearinghouse registry and ineligible for government procurement.

6.3 Phase 3 (Year 3–5): Infrastructure investment

Implication 4 (National Neuro-Ethics Clearinghouse — Ministry of Technology and Science): Establish a National Neuro-Ethics Clearinghouse to vet all neuro-educational technology before use in Zambian schools. The Clearinghouse must: maintain a publicly accessible registry of certified products; require independent peer-reviewed efficacy evidence under Mandate 3.6 before certification; provide technical assistance to schools; include youth and elder representation on its advisory board; and maintain a community-provider conflict-resolution mechanism. Estimated cost: USD 150,000–250,000 to establish; USD 60,000–100,000 per year to operate.

Implication 5 (Sovereign edge processing infrastructure — Ministry of Technology and Science): Fund Architecture B (sovereign edge processing) infrastructure on a phased basis, managed by the Ministry through a Service Level Agreement with the University of Zambia Computer Science Department using MNE-Python or EEGLAB open-source pipelines. Year 3: urban flagship schools (pilot cohort of 20 schools). Year 4: peri-urban

schools. Year 5: rural schools with additional connectivity investment. Estimated cost: USD 8,000–15,000 per school; USD 3,000–8,000 additional per rural site. Note: Architecture B deployment is conditional on at least one product having met the Mandate 3.6 efficacy threshold.

Implication 11 (Continental Neuro-Rights Protocol — African Union): Support the development of a Continental Neuro-Rights Protocol under the AU's AI Strategy [2], drawing on the Neural Chilimba framework's school-level governance principles as one possible model among several. The Protocol should be developed through multi-stakeholder consultation across African jurisdictions and should accommodate distinct national philosophical and legal traditions.

Implication 12 (Infrastructure-first development funding — international development partners): International donors and development partners should make sovereign data infrastructure, neuro-forensic training capacity, and community governance capability mandatory components of any educational technology initiative in LMIC contexts. Donors providing educational neurotechnology devices to Zambian schools must simultaneously fund Data Chilimba Committee establishment and Community Glossary co-authorship as non-negotiable conditions of the grant.

7. Discussion

7.1 Theoretical contributions

This article makes three theoretical contributions to the neuroethics literature. First, it introduces neuro-extractivism as a regulatory intensification of data colonialism — not a categorically separate phenomenon, but one with three aggravating features (irreversibility-plus-compounding-inference, involuntary disclosure, and developmental site vulnerability) that require governance provisions not yet generated by existing frameworks. The distinctiveness of neuro-extractivism lies not in any single property but in the combination of all three aggravating features in a single data type collected in a single institutional context.

Second, the article grounds its governance argument in a procedural relational claim based on the epistemic content of the data rather than in the metaphysics of personhood. The argument proceeds: EEG classroom data is inherently relational in what it measures (brain-to-brain synchrony literature [25,40]); individual consent cannot adequately govern data whose epistemic content is irreducibly social; therefore, community governance is procedurally required. This grounds the governance claim in the nature of the data itself rather than in the causal influence of the community on cognitive development, which does not in itself entail governance authority over the data measuring it.

Third, the article demonstrates that the CARE Principles for Indigenous Data Governance can be operationalised within LMIC educational policy through a purpose-built institutional mechanism — the Data Chilimba Committee — providing one model for extending international indigenous data governance frameworks beyond research contexts into educational technology deployment. The Neural Chilimba is not presented as a rival to the CARE Principles but as an operationalisation of them within a specific institutional and cultural context.

The framework's generalisability to other African nations requires contextual adaptation. Ghana's sankofa, Kenya's harambee, and Ethiopia's iqub offer analogous institutional resources, but the philosophical grounding must be drawn from each community's own tradition rather than transplanted from Zambia. The Neural Chilimba is not an African governance template; it is a Zambian governance architecture that demonstrates the method of deriving governance logic from indigenous institutional traditions.

7.2 Limitations and future research

This article has six principal limitations, of which the first four were acknowledged from the outset and the final two concern the boundaries of the current research design. Limitations that persist. First, the framework has not been empirically piloted. Future research should implement the Data Chilimba Committee in real school settings and measure community trust levels, approval rates, co-optation attempts, and benefit-sharing outcomes. Second, children's voices are absent from this analysis. Future research should employ child-centred methodologies — consistent with CRC Article 12 and the child participation literature [16,43] — to understand how Zambian learners themselves conceptualise cognitive monitoring and communal data governance. Third, the article does not include a full cost-feasibility analysis; the scoping estimates in Section 6 are indicative rather than definitive. Fourth, the framework is Zambia-specific; comparative research across African nations would identify generalisable elements and inform the Continental Neuro-Rights Protocol.

Fifth, stakeholder consultations with Zambian communities, legal practitioners, and policymakers were not conducted for this article. This is the most significant methodological gap: the framework proposes governance arrangements for communities whose voices are not represented in its development. The Community Glossary co-authorship requirement within the framework is partly responsive, but it operates within a deployment context that the community did not participate in designing. Stakeholder co-design is the priority for the next research phase.

Sixth, the literature search was conducted in English; Ubuntu and Chilimba scholarship in indigenous languages was not retrieved. A supplementary French-language search of AJOL was conducted, identifying three Francophone African data governance sources incorporated where relevant; indigenous-language scholarship remains a gap.

7.3 Implementation realism: minimum viable pathway

The phased implications in Section 6 represent the full governance architecture of the Neural Chilimba. It is important to distinguish this from what is actually achievable in the near term, given Zambia's current institutional capacity. The framework's minimum viable pathway — what could realistically be implemented first, under present constraints, comprises the following: Minimum viable now: statutory clarification that neural data falls within a protected category under Section 14 (Implication 1); local-language consent templates at Grade 7 reading level (Implication 2); a procurement prohibition on educational technology that does not disclose third-party sharing or Zambian law compliance (Implication 6); a vendor efficacy disclosure requirement as a condition of school procurement approval (Mandate 3.6); and a school-level moratorium on cloud-dependent neurotechnology without auditable processing (Mandate 3.1, Architecture C restriction). These five measures require no new infrastructure and modest budgetary investment. Medium-term institutional build: DPC staff secondment and neuro-forensics unit (Implication 3); national model benefit-sharing contracts and Data Chilimba Committee templates (Implication 8); teacher training modules (Implication 7); and the National Neuro-Ethics Clearinghouse (Implication 4). Aspirational end-state: sovereign edge processing infrastructure (Implication 5); full continental protocol development (Implication 11); and community-piloted governance templates validated through stakeholder co-design.

Zambia may not currently possess the technical, fiscal, or regulatory capacity to implement the entire model simultaneously. Some elements of the framework are intentionally threshold-setting rather than immediately executable. In the near term, the framework's most realistic function is to define non-negotiable conditions of entry rather than to operationalise the full institutional stack at once.

7.4 Next research phase: co-design and field validation

The framework proposed here should be treated as Version 1 of an anticipatory governance model whose ethical adequacy, not merely its empirical robustness, depends on subsequent co-design with the communities whose normative traditions it invokes. Community co-design is not optional downstream embellishment; it is a condition of normative legitimacy for any policy translation or pilot deployment of this framework. A governance model that invokes Ubuntu's communal traditions while having been developed without participatory dialogue with those communities is internally inconsistent in a way that only co-design can resolve. The next research phase must include: structured interviews with Zambian parents, teachers, school administrators, Ministry officials, and DPC regulators; child-centred participatory workshops using age-appropriate methods consistent with CRC Article 12 (right to be heard); consultation with Zambian legal scholars and indigenous language experts on the neural data addendum and Community Glossary co-authorship; pilot testing of Data Chilimba Committee governance templates in two to three contrasting school settings (urban flagship, peri-urban, and rural); and usability testing of local-language consent materials with parents of varying literacy levels. Absent this co-design phase, the framework is a rigorous academic governance proposal rather than a community-legitimate governance instrument. We do not obscure that distinction, and we do not regard it as resolvable by intent alone.

8. Conclusion

The Chilimba works because community members trust that their contributions will be reciprocated, that the pool will be governed fairly, and that dividends will flow to those in need. This trust is not naïve — it is institutionalised through the rotating custodian, the transparent accounting, and the communal deliberation that precedes each distribution. The Neural Chilimba extends this institutional logic to the domain of children's developing minds.

The framework proposed in this article is grounded in analysis of what Zambia's legal framework currently does and does not do; what consumer neurotechnology products currently offer and withhold; what Ubuntu philosophy specifically claims about the relational constitution of personhood and its implications for data governance; and what the Chilimba tradition specifically provides in terms of institutional logic for pooling, governance, and benefit distribution. It is explicitly positioned within the bounds of moderate communitarianism, preserving individual rights within communal governance rather than displacing them.

The framework has significant limitations: it is untested, it lacks stakeholder co-design, it has no cost-feasibility analysis, and it depends on DPC enforcement capacity that does not yet exist. These limitations are acknowledged in full rather than minimised. The contribution is not a finished governance system but a conceptual and regulatory architecture that demonstrates the method, provides a legislative basis, and specifies the institutional mechanism through which community-level neuro-rights governance could function in a Zambian school setting.

The single most actionable next step is the legislative amendment described in Implication 1: a Section 14 addendum classifying neural data as a sensitive special category under Zambia's Data Protection Act. This requires no new infrastructure, no new budget, and no new institution. It requires only a decision by the DPC to propose, and the National Assembly to pass, a four-clause amendment to existing law. Without this foundation, every other element of the framework — the Committee, the benefit-sharing architecture, the processing mandates — remains voluntary and structurally vulnerable to the superior bargaining power of multinational technology corporations. With it, Zambia would become the first African nation with an explicit legislative basis for governing children's neural data in educational settings.

That is worth doing. And the mechanism to begin already exists. The Neural Chilimba is best understood not as a finished model but as Version 1 of a candidate governance framework — contextually grounded, philosophically defended, and institutionally specified, but not yet community-validated. It defines minimum conditions of legitimacy for any future educational neurotechnology deployment in Zambia and invites empirical refinement through

stakeholder co-design. The framework is intentionally selective rather than prohibitive: it does not reject educational neurotechnology in principle, but conditions its legitimacy on demonstrable efficacy, auditable data architectures, and rights-compatible governance. The child remains at the centre — as a rights-bearing subject whose protections no communal body may waive, and whose cognitive development is precisely what this governance architecture exists to safeguard.

References

1. Abebe R, Aruleba K, Birhane A, Kingsley S, Obaido G, Remy SL, et al. Narratives and Counternarratives on Data Sharing in Africa. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, New York, NY, USA: Association for Computing Machinery; 2021, p. 329–41. https://doi.org/10.1145/3442188.3445897.
2. African Union. AU Data Policy Framework | African Union 2022. https://au.int/en/documents/20220728/au-data-policy-framework (accessed June 21, 2026).
3. Shiwlani A, Hasan SU, Kumar S. Artificial Intelligence in Neuroeducation: A Systematic Review of AI Applications Aligned with Neuroscience Principles for Optimizing Learning Strategies. Journal of Development and Social Sciences 2024;5:578–93. https://doi.org/10.47205/jdss.2024(5-IV)50.
4. Haston S, Gill S, Twentyman K, Green E, Agbeleye O, Eastaugh C, et al. A Horizon Scan of Neurotechnology Innovations. Int J Environ Res Public Health 2025;22:811. https://doi.org/10.3390/ijerph22050811.
5. Schulz D, Lillo-Navarro C, Slors M, Hrabéczy A, Reuter M. Understanding societal challenges: a NeurotechEU perspective. Front Neurosci 2024;18. https://doi.org/10.3389/fnins.2024.1330470.
6. Williamson B, Pykett J, Kotouza D. Learning brains: educational neuroscience, neurotechnology and neuropedagogy. Pedagogy, Culture & Society 2026;34:515–35. https://doi.org/10.1080/14681366.2025.2521458.
7. Musole E. Governing Neural Data in the Global South: Toward a Neuro-Responsible Business Framework for Zambia. International Journal of Advanced Business Studies 2026;5:47–61. https://doi.org/10.59857/g3bvjx74.
8. Williamson B. Making markets through digital platforms: Pearson, edu-business, and the (e)valuation of higher education. Critical Studies in Education 2021;62:50–66. https://doi.org/10.1080/17508487.2020.1737556.
9. Couldry N, Mejias UA. Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject. Television & New Media 2019;20:336–49. https://doi.org/10.1177/1527476418796632.
10. Bublitz J-C. My Mind Is Mine!? Cognitive Liberty as a Legal Concept. In: Hildt E, Franke AG, editors. Cognitive Enhancement, vol. 1, Dordrecht: Springer Netherlands; 2013, p. 233–64. https://doi.org/10.1007/978-94-007-6253-4_19.
11. Brey PAE. Anticipatory Ethics for Emerging Technologies. NanoEthics 2012;6:1–13. https://doi.org/10.1007/s11569-012-0141-7.
12. Dwork C, McSherry F, Nissim K, Smith A. Calibrating Noise to Sensitivity in Private Data Analysis. In: Halevi S, Rabin T, editors. Theory of Cryptography, Berlin, Heidelberg: Springer; 2006, p. 265–84. https://doi.org/10.1007/11681878_14.
13. Fairclough S, Gilleade K. Construction of the biocybernetic loop: a case study. Proceedings of the 14th ACM international conference on Multimodal interaction, New York, NY, USA: Association for Computing Machinery; 2012, p. 571–8. https://doi.org/10.1145/2388676.2388797.
14. Williamson B. Big Data in Education: The digital future of learning, policy and practice. 1 Oliver's Yard, 55 City Road London EC1Y 1SP: SAGE Publications Ltd; 2017. https://doi.org/10.4135/9781529714920.
15. Williams KE, Burr T, L'Estrange L, Walsh K. Early childhood educators' use of neuroscience: Knowledge, attitudes, self-efficacy and professional learning. Trends Neurosci Educ 2025;38:100247. https://doi.org/10.1016/j.tine.2025.100247.
16. Lupton D, Williamson B. The datafied child: The dataveillance of children and implications for their rights. New Media & Society 2017;19:780–94. https://doi.org/10.1177/1461444816686328.
17. Prinsloo P, Slade S. Student Data Privacy and Institutional Accountability in an Age of Surveillance. In: Menon ME, Terkla DG, Gibbs P, editors. Using Data to Improve Higher Education, Rotterdam: SensePublishers; 2014, p. 197–214. https://doi.org/10.1007/978-94-6209-794-0_12.
18. Fairclough SH. Fundamentals of physiological computing. Interact Comput 2009;21:133–45. https://doi.org/10.1016/j.intcom.2008.10.011.
19. Van Den Brink R, Chavas J. The Microeconomics of an Indigenous African Institution: The Rotating Savings and Credit Association. Economic Development and Cultural Change 1997;45:745–72. https://doi.org/10.1086/452306.
20. Zuboff S. The age of surveillance capitalism: the fight for a human future at the new frontier of power. London: Profile books; 2019.
21. Fritz Z. Anticipatory regulation, anticipatory ethics: preparing for the future. Journal of Medical Ethics 2024;50:361–2. https://doi.org/10.1136/jme-2024-110131.
22. OECD. Framework for Anticipatory Governance of Emerging Technologies. OECD Science, Technology and Industry Policy Papers 2024. https://doi.org/10.1787/0248ead5-en.
23. Bhalla N, Richardson K. A Method for Analysing Ethical Issues in Emerging Technologies – the TechEthos Approach (TEAeM) n.d.
24. Howard-Jones PA. Neuroscience and education: myths and messages. Nat Rev Neurosci 2014;15:817–24. https://doi.org/10.1038/nrn3817.
25. Dikker S, Wan L, Davidesco I, Kaggen L, Oostrik M, McClintock J, et al. Brain-to-Brain Synchrony Tracks Real-World Dynamic Group Interactions in the Classroom. Curr Biol 2017;27:1375–80. https://doi.org/10.1016/j.cub.2017.04.002.

26. Kwet M. Digital colonialism: US empire and the new imperialism in the Global South. Race & Class 2019;60:3–26. https://doi.org/10.1177/0306396818823172.
27. Nosek B. Open Science Collaboration. Estimating the reproducibility of psychological science. Science 2015;349:aac4716. https://doi.org/10.1126/science.aac4716.
28. Wexler A, Thibault R. Mind-Reading or Misleading? Assessing Direct-to-Consumer Electroencephalography (EEG) Devices Marketed for Wellness and Their Ethical and Regulatory Implications. J Cogn Enhanc 2019;3:131–7. https://doi.org/10.1007/s41465-018-0091-2.
29. Pasquinelli E. Neuromyths: Why Do They Exist and Persist? Mind, Brain, and Education 2012;6:89–96. https://doi.org/10.1111/j.1751-228X.2012.01141.x.
30. Henrich J, Heine SJ, Norenzayan A. Most people are not WEIRD. Nature 2010;466:29. https://doi.org/10.1038/466029a.
31. Musole E. Cognitive Extraction and the Neuro-Chilimba Framework: An Indigenous ROSCA-Based Governance Architecture for Neural Data in Zambia. International Journal of Applied Research in Business and Management 2026;7. https://doi.org/10.51137/wrp.ijarbm.640.
32. O'Reilly C. Urban Women's Informal Savings and Credit Systems in Zambia. Development in Practice 1996;6:165–9.
33. OECD. Responsible innovation. OECD 2019. https://www.oecd.org/en/topics/responsible-innovation.html (accessed June 21, 2026).
34. Oloyede R, Adaramola V, Buza M. The Year of the Teeth: Data Protection in Africa Roundup, 2025, Projections for 2026. https://www.techhiveadvisory.africa/report/the-year-of-the-teeth-data-protection-in-africa-roundup-2025-projections-for-2026 (accessed June 21, 2026).
35. Republic of Zambia. The Data Protection Act, 2021 | National Assembly of Zambia 2021. https://www.parliament.gov.zm/node/8853 (accessed June 21, 2026).
36. Oyěwùmí O. The Invention of Women: Making an African Sense of Western Gender Discourses. NED-New edition. University of Minnesota Press; 1997.
37. Carroll SR, Garba I, Figueroa-Rodríguez OL, Holbrook J, Lovett R, Materechera S, et al. The CARE Principles for Indigenous Data Governance. Data Science Journal 2020;19:43. https://doi.org/10.5334/dsj-2020-043.
38. Mbiti JS. African religions & philosophy. Second revised and enlarged edition. Oxford: Heinemann; 2010.
39. Vygotsky LS. Mind in Society: Development of Higher Psychological Processes. Harvard University Press; 1978. https://doi.org/10.2307/j.ctvjf9vz4.
40. Bevilacqua D, Davidesco I, Wan L, Chaloner K, Rowland J, Ding M, et al. Brain-to-Brain Synchrony and Learning Outcomes Vary by Student-Teacher Dynamics: Evidence from a Real-world Classroom Electroencephalography Study. J Cogn Neurosci 2019;31:401–11. https://doi.org/10.1162/jocn_a_01274.
41. Viljoen S. A Relational Theory of Data Governance 2020. https://doi.org/10.2139/ssrn.3727562.
42. Yuste R, Goering S, Arcas BA y, Bi G, Carmena JM, Carter A, et al. Four ethical priorities for neurotechnologies and AI. Nature 2017;551:159–63. https://doi.org/10.1038/551159a.
43. United Nations (UN). Convention on the Rights of the Child. OHCHR n.d. https://www.ohchr.org/en/instruments-mechanisms/instruments/convention-rights-child (accessed June 27, 2026).
44. Farahany NA. The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology | Department of Philosophy 2023. https://philosophy.duke.edu/books/battle-your-brain-defending-right-think-freely-age-neurotechnology (accessed June 27, 2026).
45. Boulaga FE. Kwame Gyekye, An Essay on African Philosophical Thought: the Akan conceptual scheme. Cambridge: Cambridge University Press, 1987, xv + 246 pp., £27.50, ($39.50) ISBN 0 521 32525 0. Africa 1990;60:155–8. https://doi.org/10.2307/1160437.
46. Willett C. Family Matters: Feminist Concepts in African Philosophy of Culture. By Nkiru Uwechia Nzegwu. Albany, N.Y.: SUNY, 2006. Hypatia 2008;23:224–6. https://doi.org/10.1017/S0887536700018353.
47. Ienca M, Andorno R. Towards new human rights in the age of neuroscience and neurotechnology. Life Sci Soc Policy 2017;13:5. https://doi.org/10.1186/s40504-017-0050-1.
48. Agency ZS, MoH/Zambia M of H-, ICF. Zambia Demographic and Health Survey 2024 2025.
49. Ministry of Finance and National Planning. Eighth National Development Plan (8NDP) 2022-2026 | Embassy of the Republic of Zambia in Washington, D.C. n.d. https://www.zambiaembassy.org/document/eighth-national-development-plan-8ndp-2022-2026 (accessed June 28, 2026).
50. Ministry of Technology and Science (MoTS). Ministry of Technology and Science – MOTs n.d. https://www.mots.gov.zm/ (accessed June 28, 2026).
51. Twikirize JM, Mugumbate R, Nwanna CR, Masuku MM, Tusasiirwe S, Chigangaidze RK, et al. Facilitators and barriers to the use of Ubuntu values in social work and social development. Journal of Religion & Spirituality in Social Work: Social Thought 2026;0:1–29. https://doi.org/10.1080/15426432.2026.2635374.
52. Anning-Dorson T. Digital platform sovereignty through Ubuntu: Re-imagining AI-enabled work future in Africa. Electron Markets 2025;36:2. https://doi.org/10.1007/s12525-025-00864-x.
53. Ardener S. The Comparative Study of Rotating Credit Associations. The Journal of the Royal Anthropological Institute of Great Britain and Ireland 1964;94:201–29. https://doi.org/10.2307/2844382.

54. Blakemore S-J, Frith U. The learning brain: lessons for education: a précis. Dev Sci 2005;8:459–65. https://doi.org/10.1111/j.1467-7687.2005.00434.x.
55. Bublitz C, Ligthart S. The new regulation of non-medical neurotechnologies in the European Union: overview and reflection. J Law Biosci 2024;11:lsae021. https://doi.org/10.1093/jlb/lsae021.
56. Cannard C, Wahbeh H, Delorme A. Validating the wearable MUSE headset for EEG spectral analysis and Frontal Alpha Asymmetry. 2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), 2021, p. 3603–10. https://doi.org/10.1109/BIBM52615.2021.9669778.
57. Carlini N, Tramer F, Wallace E, Jagielski M, Herbert-Voss A, Lee K, et al. Extracting Training Data from Large Language Models 2021. https://doi.org/10.48550/arXiv.2012.07805.
58. Chigwada J, Ngulube P. Will artificial intelligence replace human librarians: views from librarians in a developing country? South African Journal of Libraries and Information Science 2026;92:1–11. https://doi.org/10.7553/92-2-2505.
59. Gkintoni E, Antonopoulou H, Sortwell A, Halkiopoulos C. Challenging Cognitive Load Theory: The Role of Educational Neuroscience and Artificial Intelligence in Redefining Learning Efficacy. Brain Sci 2025;15:203. https://doi.org/10.3390/brainsci15020203.
60. Haynes J-D, Sakai K, Rees G, Gilbert S, Frith C, Passingham RE. Reading hidden intentions in the human brain. Curr Biol 2007;17:323–8. https://doi.org/10.1016/j.cub.2006.11.072.
61. Immordino-Yang MH, Damasio A. We Feel, Therefore We Learn: The Relevance of Affective and Social Neuroscience to Education. Mind, Brain, and Education 2007;1:3–10. https://doi.org/10.1111/j.1751-228X.2007.00004.x.
62. Libet B. Unconscious cerebral initiative and the role of conscious will in voluntary action. Behavioral and Brain Sciences 1985;8:529–39. https://doi.org/10.1017/S0140525X00044903.
63. Metz T. Toward an African Moral Theory. Journal of Political Philosophy 2007;15:321–41. https://doi.org/10.1111/j.1467-9760.2007.00280.x.
64. Musole E. Neuro-Ubuntu Ma□at Ethics for Neuro-AI Governance in Africa. International Journal of Advanced Business Studies 2026;5:104–26. https://doi.org/10.59857/mp49zh15.
65. Nissenbaum H, editor. Privacy in context: technology, policy, and the integrity of social life. Stanford, Calif: Stanford Law Books; 2010.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- OnlineResource1GlossaryEMSK.docx
- SDFile1ProductDocumentation.pdf
- SDFile2RubricandScoring.pdf