



Full length article

Neuro-Ubuntu Ma'at Ethics for Neuro-AI Governance in Africa

Edward Musole

University of New England, Australia

ORCID: <https://orcid.org/0009-0002-6670-3078>

*Corresponding e-mail: edwardmusoleresearch@gmail.com

Article Info

Received: 15.03.2026

Accepted: 05.05.2026

Available online: 30.05.2026

Keywords:

Neuroethics; Ubuntu; Ma'at; Neuro-Ubuntu Ma'at Ethics; neuro-AI governance; neural data governance; cognitive extractivism; African business ethics; comparative neurotechnology governance

DOI:

<https://doi.org/10.59857/mp49zh15>

ABSTRACT

This article develops a normative framework for governing neuro-AI in African business and policy contexts. Drawing on a structured literature review of neuroethics, neural data governance, neurorights, and African philosophy, the article argues that prevailing regulatory approaches remain insufficient for many African settings because they rely predominantly on individualist models of autonomy and consent. In response, the article proposes Neuro-Ubuntu Ma'at (NUM) Ethics, a synthetic framework integrating Ubuntu, Ma'at, and broader traditions of African ethical thought. The framework is organised around five principles—mind-dignity, relational personhood, Ma'atian balance, character-centred governance, and communal responsibility—and translated into governance mechanisms relevant to firms, regulators, and civil society. The article also develops the concept of cognitive extractivism to describe the extraction, commodification, and monetisation of neural data without adequate protections for dignity, mental privacy, or fair benefit-sharing. Situating NUM Ethics within a four-paradigm global typology of neurotechnology governance—rights-based, market-based, state-centric, and relational-communitarian—the article demonstrates the framework's comparative advantages and grounds it in existing African data protection and neuroethics developments. The article contributes to business ethics and technology governance by articulating an African-centred approach to neurotechnology regulation that is philosophically grounded, comparatively situated, and institutionally actionable.

1. Introduction

Advances in neurotechnology are rapidly reshaping the ethical and regulatory landscape of contemporary business. Brain-computer interfaces, cognitive monitoring systems, and neuromarketing tools are increasingly moving from experimental environments into commercial practice. These developments raise complex questions concerning mental privacy, autonomy, labour governance, and the commodification of neural data (Farahany, 2023; Ienca & Andorno, 2017; UNESCO, 2023). Recent regulatory and market developments, including the commercial approval of at least one implantable brain-computer interface device in China, indicate that neurotechnology is no longer a speculative frontier. It is becoming an organised market whose standards, deployment logics, and commercial incentives will increasingly affect jurisdictions that did not originate the technology (Cheung, 2026; Reuters, 2026; Xinhua, 2026).

For African states and firms, this transition is particularly consequential. Much of the continent remains positioned within global technology markets as an importer rather than a norm-setting producer. At the same time, regulatory systems across Africa are developing unevenly, with data protection regimes expanding in scope but often remaining silent on the specific risks associated with neural data (Eke, 2024; Musole, 2026). Existing legal instruments in many African jurisdictions do not yet classify neural data as a distinct sensitive category, nor do they comprehensively address cognitive liberty, collective data interests, or the governance implications of workplace neuro-surveillance (Botes et al., 2025; Eke, 2024; Musole, 2026). More fundamentally, many regulatory models remain strongly shaped by liberal individualist assumptions that may not adequately reflect African understandings of personhood, relationality, and communal responsibility (Andanda & Düwell, 2024; Matshabane et al., 2024).

This article argues that African neuro-AI governance requires a framework that is not merely derivative of external models but grounded in the continent's own philosophical resources and institutional realities. The central claim is that effective governance in African business and policy contexts requires a synthetic paradigm combining relational personhood, truth-oriented justice, balance, dignity, institutional character, and communal accountability (Karenga, 2004; Metz, 2022; Molefe, 2019; Obenga, 2004). To this end, the article develops Neuro-Ubuntu Ma'at (NUM) Ethics.

The article makes three contributions. First, it extends neurorights scholarship by incorporating a relational account of personhood derived from African philosophy, thereby challenging the predominantly individualist assumptions that underpin much contemporary neuroethics (Ienca & Andorno, 2017; Metz, 2022; Wiredu, 1996). Second, it develops the concept of cognitive extractivism as a targeted extension of data colonialism debates into the domain of neural data, particularly in Global South contexts where governance asymmetries may enable the commodification of cognitive signals without adequate safeguards (Couldry & Mejias, 2019; De Sousa Santos, 2018; Ndlovu-Gatsheni, 2018). Third, it introduces NUM Ethics as an integrative framework that translates African ethical traditions into institutional governance mechanisms applicable to firms, regulators, and civil society actors.

To structure the inquiry, the article addresses four research questions. What trajectories in contemporary neuroethics scholarship are most relevant to African business and governance contexts? How do emerging African approaches to consent, relational ethics, and technology governance revise prevailing frameworks? How might traditions such as Ubuntu and Ma'at be synthesised into a coherent normative paradigm for neuro-AI governance? What institutional mechanisms could translate such a paradigm into practical governance tools for firms, regulators, and communities?

A distinct contribution of this article is its explicit articulation of the theoretical architecture underpinning NUM Ethics. The framework is best understood as a middle-range governance theory: it bridges normative African ethics, human rights law, and institutional design rather than remaining at the level of philosophical abstraction. Its purpose is to convert normative principles into governance consequences.

NUM Ethics is organised around two mutually reinforcing theoretical pillars. The first pillar is Relational Personhood Theory, drawn from Ubuntu, Akan ethics, and Metz's (2022) relational moral theory. This pillar holds that persons are not isolated atoms but are socially constituted or at least socially situated in morally significant ways. Agency, dignity, and vulnerability are shaped by relationships, and governance must therefore

attend to communal as well as individual interests. The second pillar is Cosmological Balance Theory, derived from Ma'at, which foregrounds truth, justice, balance, harmony, and reciprocity as ordering principles for ethical governance (Karenga, 2004; Obenga, 2004). Where Relational Personhood asks who is affected and how relationships matter, Cosmological Balance asks whether governance sustains moral equilibrium across competing values such as innovation, safety, dignity, and fairness.

These two pillars are supplemented by three further conceptual resources. Neurorights theory (Ienca & Andorno, 2017; Yuste et al., 2017) supplies the risk vocabulary of mental privacy, cognitive liberty, agency, and psychological continuity. Indigenous African business ethics (Gyekye, 1995; Molefe, 2019) contributes a model of institutional character, stewardship, and moral formation. Cognitive extractivism (Couldry & Mejias, 2019; De Sousa Santos, 2018; Ndlovu-Gatsheni, 2018) identifies the political economy problem the framework is designed to resist: the asymmetrical extraction of neural value under conditions of weak bargaining power and uneven regulatory capacity.

Taken together, these five conceptual resources generate the analytical sequence of the article: (1) problem identification—neural data poses distinctive governance risks not adequately addressed by conventional data protection or individualist neurorights; (2) gap diagnosis—existing frameworks lack relational, communal, and political-economy dimensions essential for African contexts; (3) normative intervention—NUM Ethics synthesises the five resources into a coherent paradigm; and (4) institutional output—the paradigm is translated into four governance tools. Figure 1 (NUM Governance Architecture) visually summarises this flow from theoretical pillars to principles to institutional tools to ethical outcomes. The paper's sections follow this sequence: literature review (Section 2), methodology (Section 3), results and analysis including African governance landscape (Section 4), theoretical contribution (Section 5), comparative perspective (Section 6), the NUM framework and its operationalisation (Sections 7–8), policy application and discussion (Sections 9–11), and limitations and conclusion (Sections 12–13).

2. Literature Review

2.1 Neuroethics, neural data, and the changing scope of governance

Neuroethics emerged in response to the ethical, legal, and social implications of advances in neuroscience and neurotechnology. Its classic concerns include mental privacy, personal identity, agency, enhancement, and the moral implications of intervening in the brain (Beauchamp & Childress, 2019; Ienca & Andorno, 2017). As neurotechnology has moved closer to everyday application, the field has expanded beyond clinical and philosophical inquiry toward questions of business ethics, regulation, data governance, and political economy (Farahany, 2023; UNESCO, 2023).

A major theme in this literature is that neural data is not merely another category of personal information. For governance purposes, it is useful to distinguish three levels: raw neural signals captured by devices such as electroencephalography (EEG) headsets or implantable electrodes; processed neural data produced by algorithmic interpretation of those signals; and inferred cognitive profiles—predictive constructs about thought, preference, or vulnerability generated from processed data. This distinction matters because different governance tools may be appropriate at each level, and regulatory frameworks that focus only on raw data processing may fail to address the harms generated at the inference layer. Neural signals in any form are derived directly from the substrate of cognition and may reveal intention, emotion, attention, fatigue, preference, or

vulnerability in ways that are unusually intimate and difficult to contest (Farahany, 2023; OECD, 2019). Unlike many behavioural datasets, neural data may disclose information that individuals have neither articulated nor consciously recognised. That feature generates distinctive concerns regarding mental privacy, autonomy, manipulation, and downstream commercial inference. For this reason, scholars increasingly argue that conventional data protection models are insufficient for neurotechnology governance (Farahany, 2023; OECD, 2019; UNESCO, 2023).

2.2 The normative promise and limits of neurorights

The neurorights movement represents one of the most important recent developments in neuroethics. It advances claims around cognitive liberty, mental privacy, mental integrity, and psychological continuity as a means of updating human rights discourse for the age of neurotechnology (Ienca & Andorno, 2017; Yuste et al., 2017). These concepts are valuable because they identify harms that are not always well captured by conventional privacy or consumer protection law. They also provide a vocabulary through which legislators and courts can begin to address brain-data harms.

Yet the literature also reveals important limitations. Neurorights scholarship often presumes an ethical subject imagined primarily as an independent chooser whose core interests are expressed through consent, control, and non-interference. That model remains useful, but it is incomplete for settings in which decision-making is often relational and harms are experienced not only as private injuries but as disruptions of household stability, communal trust, and social belonging (Andanda & Düwell, 2024; Botes et al., 2025). In addition, neurorights discourse may focus heavily on rights recognition while devoting less attention to institutional capacity, benefit-sharing, and the commercial conditions under which neural data is extracted and monetised (Eke, 2024; Matshabane et al., 2024). The challenge, therefore, is not to reject neurorights but to situate them within a broader governance architecture.

2.3 Relational ethics and the limits of Western individualism

Mainstream neuroethics remains strongly shaped by liberal individualist accounts of autonomy. Informed consent is often treated as the procedural expression of an independent individual's choice, made free from external interference (Beauchamp & Childress, 2019). This model offers institutional clarity, but it assumes social conditions that are often incomplete in African settings, including high information symmetry, strong bargaining power, stable literacy, and a conception of selfhood only lightly mediated by family or community (Botes et al., 2025).

By contrast, much African ethical thought treats the person as socially constituted or at least socially situated in a morally significant sense (Metz, 2022; Molefe, 2019; Wiredu, 1996). On this view, consultation does not necessarily negate voluntariness. It may instead contribute to ethically meaningful choice. Relational autonomy can therefore strengthen rather than weaken consent validity, especially where technological literacy varies and trusted social networks improve comprehension and reduce vulnerability to coercion (Botes et al., 2025; Ewuoso & Hall, 2019). This matters for neurotechnology because the consequences of cognitive intrusion often extend beyond the individual to families, workplaces, and communities.

2.4 Ubuntu as a resource for neuro-AI governance

Ubuntu remains the most prominent African philosophical resource in contemporary ethics scholarship. Often expressed through the proposition that a person is a person through other persons, Ubuntu offers a powerful counterpoint to atomised accounts of agency and dignity (Ewuoso & Hall, 2019; Metz, 2022). Its relevance to neuroethics lies in its insistence that moral life is organised around care, reciprocity, solidarity, and mutual recognition.

In governance terms, Ubuntu has at least three implications. First, it supports dialogical and iterative consent rather than one-time transactional signatures. Second, it supports the view that communities may have legitimate interests in how group-linked or socially consequential neural data is collected, interpreted, and commercialised. Third, it positions organisations as moral actors with stewardship duties toward workers, consumers, and affected communities rather than as neutral data processors.

This article adopts a moderate relational interpretation of Ubuntu, following Metz (2022) and Ewuoso and Hall (2019), in which persons are socially situated but retain independent moral standing and the power of refusal. This contrasts with strong constitutive views—for instance, readings that treat personhood as entirely constituted by community membership—which NUM Ethics rejects as inconsistent with the protection of cognitive liberty and individual agency. It also contrasts with purely individualist critiques that treat communal consultation as a threat to voluntariness; NUM Ethics instead treats such consultation, where appropriately structured, as a resource for comprehension and legitimacy. The article further recognises that Ubuntu is internally contested, with some interpretations emphasising relational personhood more strongly and others preserving greater independent weight for personal agency (Metz, 2022; Molefe, 2019). NUM Ethics accordingly preserves individual veto while broadening the ethical context of decision-making.¹

2.5 Ma'at as a governance ethic of truth, justice, and balance

If Ubuntu supplies a relational account of personhood, Ma'at contributes a complementary account of truth, justice, balance, harmony, and right order. Dedicated scholarship on Ma'at makes clear that it is not merely a symbolic concept but a substantive moral ideal within classical African ethics (Karenga, 2004; Obenga, 2004). In the context of neuro-AI governance, Ma'at is useful because it foregrounds several requirements often obscured by procedural compliance. Truth requires candour about what neurotechnology can and cannot do. Justice requires scrutiny of who bears the burdens and who receives the benefits. Balance requires disciplined evaluation of innovation, safety, dignity, privacy, and legitimacy. Reciprocity requires that organisations deriving value from neural data return protection, accountability, and fair value to those from whom such data is drawn. Ma'at is especially valuable in business ethics because it asks not only whether conduct is formally permissible, but whether it sustains moral equilibrium. In a field marked by asymmetries of expertise, opaque inference, and strong market incentives, such a standard helps guard against hype, concealment, and extractive deployment. The article draws principally on the scholarly reconstructions of Karenga (2004) and Obenga (2004), treating

¹NUM Ethics adopts a moderate relational position: persons are socially situated and relationships shape moral agency, but individuals retain independent moral standing and the power of refusal. This contrasts with strong constitutive views in which personhood is entirely constituted by community, which NUM Ethics rejects as inconsistent with the protection of cognitive liberty and individual agency. It also contrasts with purely individualist critiques that treat communal consultation as necessarily coercive; NUM Ethics preserves individual veto absolutely while using relational consultation as a resource for comprehension and legitimacy.

Ma'at as a normative resource rather than a settled historical claim. Future research should engage more deeply with Francophone and Arabophone Egyptological scholarship to subject this interpretation to critical scrutiny.

2.6 Indigenous African business ethics and institutional character

A wider set of African ethical traditions reinforces the case for synthesis. Akan ethics foregrounds character, responsibility, and the moral significance of personhood within community (Gyekye, 1995; Molefe, 2019). More broadly, indigenous African business ethics often resist reducing ethical conduct to formal rule compliance and instead emphasise integrity, self-mastery, social obligation, dignity, and practical wisdom (Metz, 2022; Molefe, 2019).

This is consequential because neurotechnology governance will not be secured by law alone. Institutions may formally comply with disclosure rules while still operating in manipulative, extractive, or degrading ways. Organisational character matters. Managerial disposition matters. Indigenous African business ethics therefore extends neuroethics beyond rights language into the terrain of stewardship, covenant, and moral formation.

2.7 Why this synthesis?

The selection of Ubuntu, Ma'at, and broader indigenous African business ethics is not intended to suggest that African ethics is uniform or exhausted by these traditions. African ethical traditions are philosophically diverse and regionally differentiated. NUM Ethics is therefore proposed as a heuristic synthesis rather than as a claim of uniform continental doctrine. These three pillars are selected because they jointly address the principal gaps identified in the literature review. Ubuntu offers an account of relational personhood and social dignity. Ma'at offers a governance ethic of truth, justice, and equilibrium. Indigenous business ethics contributes institutional character and stewardship. Taken together, they respond to recurring concerns with autonomy, extractive value capture, organisational accountability, and the need for a governance framework capable of practical application. NUM Ethics is not proposed because existing models are without value, but because none of them individually addresses the full range of governance challenges identified here—specifically, the intersection of relational personhood, communal interests, cognitive extractivism, and institutional capacity in African contexts.

3. Methodology

This study employs a structured literature review designed to synthesise interdisciplinary scholarship on neuroethics, neural data governance, neurorights, and African moral philosophy (Pati & Lorusso, 2018; Snyder, 2019). The review is structured rather than fully systematic in the narrow PRISMA sense, although it adopts transparent search and screening procedures to support analytic credibility and replicability. A structured review is methodologically justified for this article because the goal is normative and conceptual synthesis rather than the aggregation of empirical intervention studies; prioritising analytical coherence and thematic depth over exhaustive enumeration is appropriate for frameworks of this type.

Searches were conducted across Scopus, Web of Science, and Google Scholar using combinations of terms relating to neurotechnology, brain-computer interfaces, neural data, neurorights, African ethics, Ubuntu, Ma'at, business ethics, labour governance, and technology regulation. The search process yielded approximately 312 records. After duplicate removal, 241 records remained. Title and abstract screening reduced the set to 128 records, from which 74 full-text sources were reviewed. A final corpus of 52 sources was included in the thematic synthesis.

Table 1. Structured literature review screening stages

Review stage	Number of records
Initial records identified	312
After duplicates removed	241
After title and abstract screening	128
Full-text sources reviewed	74
Final sources included in thematic synthesis	52

Sources were included where they made a substantive contribution to the ethical, legal, governance, or business dimensions of neurotechnology. Peer-reviewed journal articles, scholarly books, and authoritative policy documents were prioritised. The analysis followed Braun and Clarke's (2006) thematic coding approach, with attention to recurring concepts including autonomy, relationality, dignity, justice, privacy, cognitive liberty, consent, institutional capacity, benefit-sharing, technology transfer, and organisational responsibility. The methodology has limitations. The analysis relied principally on English-language sources, which may underrepresent Francophone, Lusophone, and Arabophone scholarship.

4. Results and Analysis

4.1 Neural data as a distinct governance category

The literature strongly supports the view that neural data should be governed as a distinct and especially sensitive category of information (Farahany, 2023; Ienca & Andorno, 2017; OECD, 2019). Neural data concerns the substrate of cognition and may disclose patterns relating to thought, preference, emotion, fatigue, intention, and identity. It may enable inferential claims that are more intimate and potentially more consequential than those associated with many conventional datasets. Once processed, combined, or commercialised, such data may generate harms that cannot be meaningfully undone.

This result has direct relevance for business governance. In commercial settings, organisations are incentivised to collect high-value data that promises prediction, segmentation, optimisation, and control. Neural data is therefore likely to attract interest not only in medicine, but also in labour management, advertising, financial assessment, education technology, and consumer analytics (Farahany, 2023; UNESCO, 2023). Without early safeguards, market incentives are likely to normalise collection before norms of restraint become institutionally established.

4.2 Four governance gaps in African contexts

A second major finding is the existence of four interrelated governance gaps across many African settings. The first is categorical silence: neural data is rarely recognised explicitly as a distinct form of sensitive information in African data protection legislation. The second is weak collective governance, whereby communities, households, or worker groups have little formal standing in decisions involving aggregated or socially consequential neural data. The third is underdeveloped cognitive liberty and mental privacy protection, particularly in labour and consumer contexts. The fourth is the absence of robust benefit-sharing and redress mechanisms, leaving open

the possibility that individuals and communities bear the burdens of cognitive data extraction while receiving little value or protection in return (Eke, 2024; Matshabane et al., 2024; Musole, 2026).

These gaps are magnified by practical constraints. Regulatory institutions may lack specialist expertise, investigative capacity, or resources for oversight. Businesses may adopt neurotechnology before sector-specific guidance exists. Workers and consumers may have limited bargaining power or limited understanding of how neural data can be repurposed. The combined effect is a governance environment in which legal ambiguity and institutional weakness can be exploited by technologically sophisticated actors.

4.3 Cognitive extractivism and workplace neuro-surveillance

One of the clearest implications of the review is that workplace neuro-surveillance constitutes a critical stress test for governance adequacy. Technologies marketed for fatigue detection, concentration monitoring, emotional assessment, or safety optimisation can quickly become instruments of intrusive surveillance. In sectors such as mining, transport, manufacturing, logistics, and security, employers may invoke legitimate safety concerns while simultaneously expanding managerial control over workers' cognitive lives (Farahany, 2023).

This article uses cognitive extractivism as a targeted extension of data colonialism analysis into neurotechnology governance. It refers to the extraction, commodification, and monetisation of neural or cognitive data from persons and communities without equitable governance, fair benefit-sharing, or adequate safeguards for mental privacy and dignity (Couldry & Mejias, 2019; De Sousa Santos, 2018; Ndlovu-Gatsheni, 2018). The concept is especially useful in African contexts because it captures how neural data can become a new site of asymmetrical value extraction under conditions of limited bargaining power and uneven regulatory capacity.

For analytical and regulatory precision, cognitive extractivism may be understood at three levels. Primary cognitive extractivism involves the direct extraction of neural signals from living persons through devices such as EEG headsets or implantable electrodes. Secondary cognitive extractivism involves the repurposing and aggregation of existing neural datasets—for instance, clinical data redirected to commercial inference without renewed consent or benefit-sharing. Tertiary cognitive extractivism involves the algorithmic generation of cognitive profiles without direct neural access, through inference from behavioural data, eye-tracking, digital exhaust, or other proxies. This third form is particularly consequential for governance because it does not require any neural device and therefore may fall outside regulatory frameworks focused exclusively on device oversight. Together, these three levels indicate that the governance challenge extends well beyond direct device use into the broader ecosystem of cognitive data inference.

In African contexts, such extractivism would not merely harm isolated individuals. It could deepen existing patterns of dependency by turning populations into sites of cognitive value capture under governance conditions they did not meaningfully shape.

4.4 Technology transfer and strategic vulnerability

A further finding concerns the strategic implications of rapid neurotechnology commercialisation in more powerful jurisdictions. When exported systems arrive with embedded assumptions about surveillance, monetisation, risk tolerance, or consent, importing jurisdictions may inherit governance models they did not

normatively author (Cheung, 2026; Reuters, 2026; Xinhua, 2026). For African states, the issue is not only technological dependency but also ethical dependency.

The literature therefore supports a proactive rather than reactive governance strategy. If standards are articulated only after imported systems have become entrenched, regulatory leverage will be reduced. A coherent African framework can strengthen procurement review, licensing decisions, labour safeguards, and cross-border technology transfer negotiations.

4.5 Emerging African Approaches to Data and Neuro-Governance

A critical contextual finding is that African governance in this domain is better characterised as data-centric, institutionally uneven, and not yet neuro-specific than as absent. This reframing matters because it positions NUM Ethics as a harmonising meta-framework extending existing African governance infrastructure rather than as a proposal drafted in a regulatory vacuum.

At the continental level, the African Union's Convention on Cyber Security and Personal Data Protection (Malabo Convention, 2014) represents the most ambitious regional instrument. The Convention entered into force in June 2023 following ratification by the requisite fifteen states. It establishes a continental framework for personal data protection, cybersecurity, and electronic transactions. While the Malabo Convention does not explicitly address neural data, its data protection provisions—requiring member states to establish data protection authorities, implement principles of purpose limitation and proportionality, and classify sensitive personal data with heightened protections—provide foundational architecture on which neuro-specific obligations could be layered. The African Union's broader digital and AI governance work, including UNESCO-aligned normative development, further signals continental momentum (UNESCO, 2025).

At the national level, a growing number of African states have enacted comprehensive data protection legislation. South Africa's Protection of Personal Information Act 4 of 2013 (POPIA) classifies biometric information as 'special personal information' subject to heightened protection, and while it does not explicitly reference neural data, its biometric and health data categories could be interpreted to encompass certain forms of neural signal. South Africa also leads continental neuroethics scholarship: Botes et al. (2025) conduct a dedicated neuroethical inquiry into consent and cognitive liberty, while Matshabane et al. (2024) advance Africa-specific neuroethics guidelines. Kenya's Data Protection Act 2019 establishes a Data Commissioner, requires data protection impact assessments, and classifies sensitive personal data including biometric data with heightened protections, but contains no neural-specific provisions. Nigeria's Nigeria Data Protection Act 2023—which replaced the Nigeria Data Protection Regulation of 2019 and established the Nigeria Data Protection Commission—similarly provides a strengthened legislative framework for sensitive data without addressing neural data as a distinct category. Rwanda's Law No 058/2021 on the Protection of Personal Data and Privacy classifies biometric and health data as sensitive. Zambia's Data Protection Act 2021 establishes a general sensitive data category without neural-specific provisions (Musole, 2026).

The common pattern across these instruments is instructive: African data protection frameworks classify biometric and health data as sensitive and subject to heightened protection, but none yet explicitly recognises neural data as a distinct category warranting specific governance obligations. Cognitive liberty, workplace neuro-surveillance, benefit-sharing from neural data, and collective data interests remain largely unaddressed. This is the precise gap that NUM Ethics is designed to fill—not by displacing existing frameworks but by providing the

neuro-specific normative extension that bridges current legislation and the emerging realities of neurotechnology deployment. It is also worth noting that several African scholars specifically cited in this article—Eke (2024), Matshabane et al. (2024), and Botes et al. (2025)—are themselves generating the neuroethics knowledge base on which future African legislative development will draw. NUM Ethics aims to operationalise and extend that emerging scholarly foundation.

Table 1B. Selected African data protection instruments and neural data governance status

Jurisdiction	Primary Instrument	Sensitive Data Coverage	Neural Data Status	Gap for NUM Ethics
South Africa	POPIA 2013	Biometric & health data	Not explicitly classified	Neural-specific category; relational consent
Kenya	Data Protection Act 2019	Biometric & health data	Not mentioned	Institutional tools (covenants, NUM councils)
Nigeria	Data Protection Act 2023	Sensitive personal data	Not explicitly classified	Character-centred governance for firms
Rwanda	Law No 058/2021	Biometric & health data	Not specified	Ma'atian balance for procurement review
Zambia	Data Protection Act 2021	General sensitive category	Not addressed	Cognitive extractivism as legal category
AU (Continental)	Malabo Convention (in force 2023)	Sensitive personal data; proportionality	Not addressed	Harmonising neuro-specific protocol

Note: Data protection instrument details verified against publicly available legal texts. Neural data status assessed against explicit legislative provisions as of 2026.

5. Theoretical Contribution

This article advances three contributions to the literature on business ethics and technology governance. First, it extends neurorights scholarship by incorporating a relational account of personhood derived from African philosophy, thereby challenging the predominantly individualist assumptions that underpin much contemporary neuroethics. Second, it develops the concept of cognitive extractivism to describe the political economy of neural data in contexts where cognitive signals may be commodified without adequate safeguards, including at primary, secondary, and tertiary levels. Third, it introduces NUM Ethics as an integrative framework that translates African ethical traditions into institutional governance mechanisms applicable to firms, regulators, and civil society actors.

The article's novelty does not lie in claiming that Ubuntu, Ma'at, or neurorights are individually new. Rather, its contribution lies in their combined synthesis into a middle-range governance theory, in the explicit comparative positioning of that theory within a global typology of neurotechnology governance models, and in the translation

of philosophical principles into practical institutional tools. The following section develops that comparative positioning.

6. Global Models of Neurotechnology Governance: A Comparative Typology

NUM Ethics is best understood not as a replacement for existing legal regimes but as a framework for governance design in contexts where relational ethics and communal interests play a significant role. To articulate its comparative position, this section develops a four-paradigm typology of global neurotechnology governance. The four paradigms are: rights-based (exemplified by European and Chilean approaches), market-based (exemplified by the United States), state-centric (exemplified by China), and relational-communitarian (proposed by NUM Ethics). This typology serves not merely as a descriptive survey but as a taxonomic contribution that situates NUM Ethics as a genuinely distinct governance paradigm rather than a regional variant of existing approaches.

6.1 The European rights-based model

The dominant logic of European neurotechnology governance is rights-based and compliance-oriented. The General Data Protection Regulation (EU) 2016/679 (GDPR) provides the primary data protection architecture. Neural data does not constitute a named category under the GDPR, but it may qualify as sensitive data under Article 9 where it constitutes health data, biometric data used for identification, or data concerning mental health. The GDPR's consent architecture requires lawful basis, purpose limitation, data minimisation, and the right to withdraw consent. Data protection impact assessments (Article 35) are required for high-risk processing, which would likely encompass most commercial neurotechnology deployments. The European Union AI Act (Regulation (EU) 2024/1689, which entered into force on 1 August 2024 and is fully applicable from 2 August 2026) establishes a risk-based classification framework for artificial intelligence systems but does not designate neurotechnology as a specific high-risk category, addressing instead the AI applications that may interpret or act upon neural data.

The GDPR's limitations for neurotechnology governance are instructive. It was designed for data processing regulation rather than protection of the cognitive substrate itself. A framework focused on lawful processing of data after collection does not directly address the moment of neural signal capture, the coercive conditions under which consent may be sought in employment contexts, or the governance of inferential cognitive profiles that may not technically constitute the personal data of identifiable individuals. Chile's 2021 constitutional amendment to protect 'mental integrity' and 'mental indemnity' as fundamental rights (Article 19(1) of the Political Constitution), and subsequently by Bulletin 13.828-19 — an implementing bill that has passed the Senate and, as of 2026, remains under consideration in the Chamber of Deputies — represent a rights-based advance on the European model because they constitutionalise the cognitive substrate rather than merely regulating downstream data processing. However, Chile's implementing legislation has been critiqued for symbolic strength combined with limited institutional enforcement mechanisms and uncertain application to commercial contexts.

The strength of rights-based models lies in strong procedural safeguards, individual empowerment, and democratic accountability. Their limitation for African governance is their over-individualised conception of the harms at stake, their weak vocabulary for communal harms and collective benefit-sharing, and their tendency to prioritise procedural compliance over substantive ethical culture in institutions.

6.2 The United States market-based model

The United States presents a contrasting governance logic: predominantly market-based, with fragmented federal and state-level regulation and a strong orientation toward private-sector development. There is no comprehensive federal privacy statute governing neural data. The sectoral patchwork of federal privacy law—including HIPAA for health information, the FTC Act’s unfair or deceptive practices standard, and limited state-level privacy statutes—was not designed for neural data and leaves significant regulatory gaps, particularly for consumer and workplace neurotechnology. Several US states have enacted biometric privacy laws (most notably the Illinois Biometric Information Privacy Act, 740 ILCS 14, 2008), and Colorado enacted neural data protections through its 2024 amendment to the Colorado Privacy Act (Colo. Rev. Stat. § 6-1-1301 et seq., as amended 2024). These are important developments, but they remain jurisdictionally fragmented and do not address the full spectrum of neurotechnology governance challenges.

The US model is directly relevant to the cognitive extractivism thesis. The dominance of private-sector neurotechnology development—including major commercial brain-computer interface programmes—combined with minimal federal constraint creates conditions under which neural data commodification can proceed with limited public accountability. This is precisely the market logic that NUM Ethics is designed to contest: the default assumption that neural data is available for commercial exploitation unless a specific legal prohibition applies.

The strength of market-based models lies in innovation dynamism and private-sector capacity. Their limitation for African governance is that they assume conditions of market competition, consumer sophistication, and litigation capacity that may not obtain in African settings, and they provide no vocabulary for communal interests or benefit-sharing obligations.

6.3 China’s state-centric model

China’s approach to neurotechnology governance is state-centric: organised around industrial strategy, national standards, and coordinated regulatory development. China identified brain-computer interfaces as ‘an industry of the future’ in its 2026 government work report (Xinhua, 2026) and commercially approved Neuracle’s implantable brain-computer interface device in March 2026—the first such regulatory approval globally for a commercial implantable BCI (Cheung, 2026; Reuters, 2026). China’s Personal Information Protection Law (2021) classifies biometric information as sensitive personal data subject to separate consent requirements, providing a relevant data protection foundation. State-level neurotechnology governance in China has also been shaped by emerging national standards on BCI safety, efficacy, and classification, distinguishing therapeutic from augmentative applications and classifying devices by invasiveness—a form of technology-specific governance classification that European frameworks have not always made explicit.

China’s state-centric model is a useful comparator not because it is normatively preferable, but because it demonstrates what governance architecture designed specifically for neurotechnology categories—rather than adapted from general data protection—can look like. Its limitations are significant, however: the model prioritises industrial development and state strategic interest over participatory communal legitimacy, provides relatively weak protections for individual cognitive autonomy relative to state interests, and generates limited accountability for extractive commercial deployment.

6.4 NUM Ethics as a fourth paradigm: relational-communitarian governance

Existing governance approaches thus operate within three dominant paradigms: rights-based, market-based, and state-centric. NUM Ethics proposes a fourth: relational-communitarian governance. In this paradigm, neuro-AI is assessed not only in terms of individual choice and procedural compliance (rights-based), market efficiency (market-based), or state industrial strategy (state-centric), but in terms of communal legitimacy, stewardship, justice, and the prevention of cognitive extractivism.

The relational-communitarian paradigm is distinctive in several ways. It treats communal interests as legally and ethically cognisable alongside individual interests. It requires governance mechanisms—brain-data covenants, NUM councils, person-in-community consent—that embed community oversight and benefit-sharing into the deployment architecture of neurotechnology, rather than treating these as voluntary CSR add-ons. And it grounds these mechanisms in African philosophical traditions—Ubuntu and Ma’at—that carry genuine normative authority within the contexts where the framework is designed to operate. This last feature is important for legitimacy: a governance framework that resonates with the values and relational structures of the communities it governs is more likely to be adopted, respected, and enforced than one imposed from outside.

Table 2. Comparative evaluation of global neurotechnology governance paradigms

Model	Core Logic	Mental Privacy	Relational Harm	Benefit-Sharing	African Context Fit	NUM Ethics Advantage
EU (GDPR/AI Act)	Rights-based; compliance	Strong (data-processing level)	Weak	Absent	Low-Medium	Adds relational personhood; communal responsibility ; Ma’atian balance
Chile (neurorights)	Rights-based; constitutional	Strong (substrate level)	Limited	Absent	Low-Medium	Adds institutional enforcement; benefit-sharing; communal governance
USA (market-based)	Market-led; fragmented	Weak federally	Absent	Absent	Very Low	Directly counters cognitive extractivism; adds stewardship obligations

China (state-centric)	State-led; industrial	Moderate	Absent	Partial (state-directed)	Low	Adds communal legitimacy; anti-extractive protections; participatory consent
African DPAs (baseline)	Rights-based; data-centric	Moderate (biometric/health)	Absent	Absent	Medium (foundational)	Extends with neural-specific category; cognitive liberty; NUM councils
NUM Ethics	Relational-communitarian	Strong (substrate + inference)	Strong	Required	High	— (proposed framework)

Note: Assessments based on published legal instruments and scholarly commentary as of 2026. EU AI Act entered into force August 2024; high-risk AI provisions fully applicable from 2 August 2026. Chile constitutional neurorights amendment: Ley 21.383 (2021); implementing bill Bulletin 13.828-19 before Chamber of Deputies as of 2026. Malabo Convention in force 8 June 2023.

6.5 Comparative advantages of NUM Ethics

The comparative analysis reveals five specific advantages of NUM Ethics over existing paradigms. First, it is more contextually legitimate in African settings because it aligns with relational personhood rather than assuming the isolated chooser that underpins rights-based and market-based models. Second, it is more protective against extractive harms because it directly addresses cognitive extractivism at primary, secondary, and tertiary levels, rather than focusing only on consent at the point of data collection. Third, it is more institutionally actionable because it translates principles into concrete governance tools—design benchmarks, consent models, brain-data covenants, and NUM councils—rather than remaining at the level of rights recognition. Fourth, it is more socially realistic because it recognises that neuro-AI harms affect families, workers, and communities as well as isolated individuals, and it provides governance mechanisms for communal interests. Fifth, it is more normatively complete because it adds truthfulness, equilibrium, stewardship, and communal responsibility to the standard rights vocabulary, addressing ethical culture and institutional virtue rather than procedural compliance alone. NUM Ethics does not claim to replace the existing models described above. In many African settings, data protection legislation aligned with GDPR-style frameworks and neurorights-informed constitutional provisions will provide the formal legal backbone of neuro-AI governance. NUM Ethics is proposed as the neuro-specific normative extension of that legal infrastructure: the layer that addresses what existing frameworks leave ungoverned, and does so in a manner grounded in African philosophical traditions and institutional realities.

7. Neuro-Ubuntu Ma'at Ethics: Conceptual Framework

NUM Ethics is proposed as a synthetic paradigm for governing neuro-AI in African business and policy landscapes. It is synthetic in two senses. Philosophically, it brings together ethical traditions too often treated in isolation through a process of hermeneutic retrieval and synthesis rather than any claim of unmediated derivation. Institutionally, it links normative reasoning to governance design. Its purpose is not to collapse African ethical thought into a single voice, but to formulate a principled framework capable of resonance across diverse settings while remaining sufficiently concrete to guide action.

The first principle is:

Mind-dignity. Because neural data concerns intimate dimensions of thought, emotion, and personhood, the mind should not be treated as an ordinary site of economic extraction. Mind-dignity imposes a presumption against intrusive collection, a requirement of strict necessity for high-impact use, and a commitment to mental privacy, mental integrity, and respect for intrinsic human worth (Farahany, 2023; Ienca & Andorno, 2017). This principle is the governance expression of the Ubuntu understanding that personhood is sacred and the Ma'atian requirement that truth and right order govern how persons are treated.

Relational personhood. Individuals remain the primary bearers of rights and must retain the power to refuse neurotechnological intervention or monitoring. At the same time, persons are embedded in relationships that shape decision-making, vulnerability, and the consequences of harm. Governance should therefore permit meaningful consultation and context-sensitive explanation without allowing communal structures to extinguish personal agency (Botes et al., 2025; Metz, 2022). This principle operationalises the moderate relational Ubuntu position described in Section 2.4.

Ma'atian balance. Ethical governance must resist both technological romanticism and reflexive prohibition. Neurotechnology should be evaluated through a disciplined balancing of innovation, utility, safety, justice, social legitimacy, and long-term moral cost. This principle also incorporates truthfulness: claims made by firms, researchers, and regulators must be candid, evidence-based, and proportionate to what the technology can actually deliver (Karenga, 2004; Obenga, 2004). Ma'atian balance provides the 'how' of governance evaluation, complementing the 'who' established by relational personhood.

Character-centred governance. Governance cannot depend solely on external rules. It also depends on the moral quality of institutions and decision-makers. NUM Ethics emphasises integrity, prudence, fairness, and stewardship as organisational virtues, drawing on the indigenous African business ethics tradition described in Section 2.6. Where institutions cultivate these qualities, compliance becomes more than procedural adherence; it becomes an expression of ethical orientation (Gyekye, 1995; Molefe, 2019). This principle distinguishes NUM Ethics from compliance-based frameworks by attending to institutional culture rather than only to rule enforcement.

Communal responsibility. The effects of neurotechnology often extend beyond isolated users to workplaces, families, social groups, and future generations. Organisations therefore bear responsibilities that are social as well as contractual. Where harm occurs, response should include repair, explanation, and institutional learning rather than narrow liability management. This principle is the governance translation of Ubuntu's insistence on care, reciprocity, and solidarity as moral foundations.

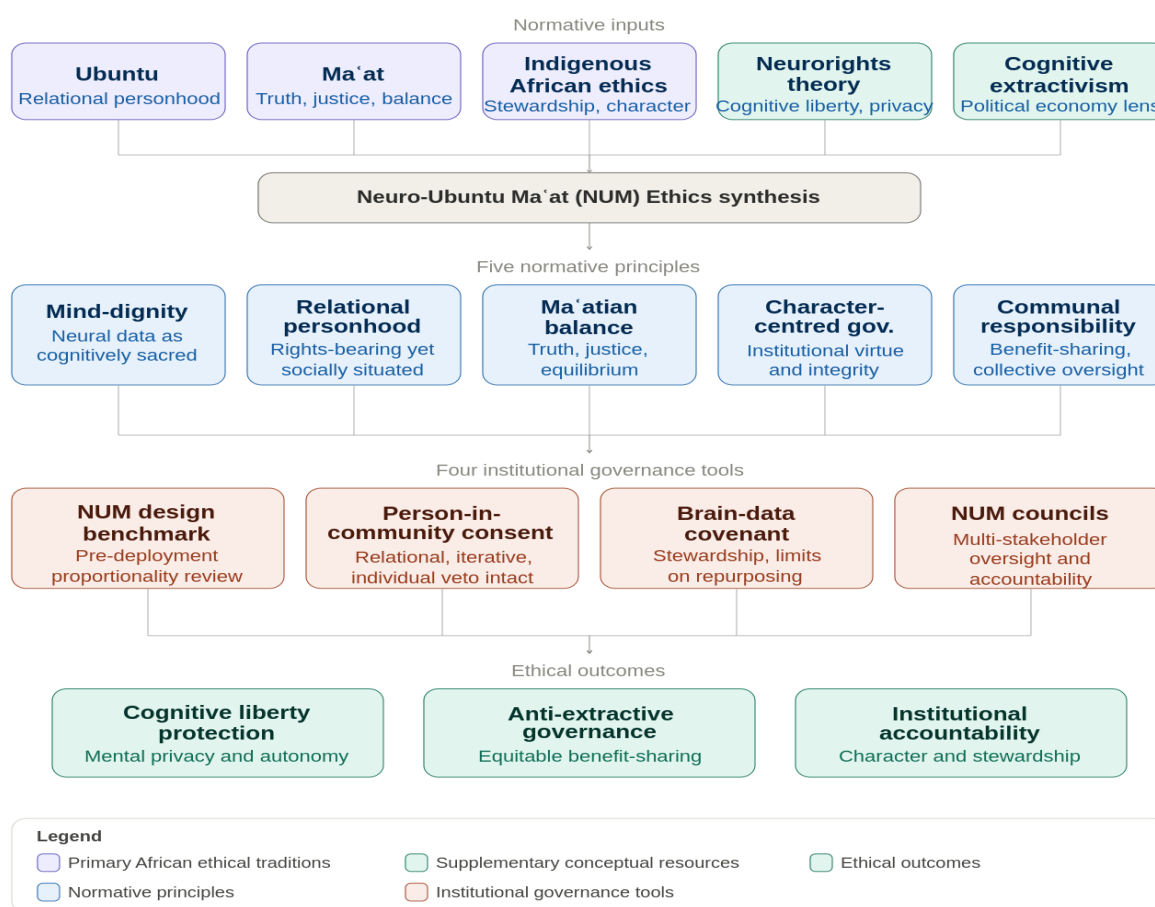


Figure 1. NUM Ethics Governance Architecture: hermeneutic synthesis linking African ethical traditions, supplementary frameworks, normative principles, institutional tools, and ethical outcomes (Musole, 2026).

Figure 1. NUM Governance Architecture

This figure illustrates a hermeneutic synthesis linking African ethical traditions (Ubuntu, Ma'at, indigenous business ethics), supplemented by neurorights and cognitive extractivism, into five normative principles (mind-dignity, relational personhood, Ma'atian balance, character-centred governance, communal responsibility), which generate four institutional governance tools (NUM design benchmark, person-in-community consent, brain-data covenant, NUM councils), directed toward three ethical outcomes (cognitive liberty protection, anti-extractive governance, institutional accountability).

Table 3 synthesises the five principles into a single reference, mapping each principle to its core normative idea and its direct governance implication. Taken together, the five principles form a coherent normative architecture: mind-dignity and relational personhood define who the framework protects and how; Ma'atian balance and character-centred governance define the evaluative and institutional standards to which actors are held; and communal responsibility defines the social and temporal scope of accountability.

Table 3. Core principles of NUM Ethics and governance implications

Principle	Core normative idea	Governance implication
Mind-dignity	The mind must not be reduced to an extractive commercial object	Heightened safeguards for neural data and a presumption against intrusive monitoring
Relational personhood	Persons are rights-bearing yet socially situated	Consent processes must preserve refusal while enabling contextual consultation
Ma'atian balance	Innovation must be governed through truth, justice, and equilibrium	Proportionality review, transparency, and anti-hype obligations
Character-centred governance	Ethical governance depends on institutional virtue as well as rules	Stewardship duties, leadership accountability, and ethical organisational culture
Communal responsibility	Neurotechnology affects communities as well as individuals	Benefit-sharing, collective oversight, and restorative redress mechanisms

The governance implications in the third column are not abstract ideals but operational directives. Each points toward a specific type of institutional practice: heightened procedural safeguards, relational consent processes, proportionality review, stewardship culture, and community oversight mechanisms. Section 8 translates these directives into the four concrete governance tools through which NUM Ethics is made institutionally actionable.

8. Operationalising NUM Ethics

For a normative framework to be practically credible, it must be institutionally translatable. NUM Ethics is therefore developed into four governance instruments. These instruments are not philosophical afterthoughts; they are the operational payoff of the normative analysis and should be understood as the framework's primary practical contribution.

The NUM design benchmark. Before any neurotechnology is deployed in a business, institutional, or regulatory setting, decision-makers should justify purpose necessity, establish proportionality, identify likely dignity and privacy harms, explain data flows, evaluate downstream commercial risks, and assess the likely effects on worker or community trust. This benchmark serves as a pre-deployment review instrument applicable to procurement decisions, vendor contracts, regulatory licensing, and internal compliance systems (Farahany, 2023; UNESCO, 2023).

Person-in-community consent. This model reconceptualises consent as an ongoing and context-sensitive process rather than a single contractual act. Information must be intelligible, jargon-free, and adapted to differing levels of digital and health literacy. Individuals may, where appropriate, invite consultation with family members, community intermediaries, or trusted advisers. The model preserves individual veto absolutely: the individual's express refusal terminates the process regardless of community preference; any consultation is advisory rather than dispositive; and refusal may not trigger punitive consequences unless an independently reviewable safety justification exists. Relational consultation is thus used to strengthen comprehension and legitimacy without displacing personal agency (Botes et al., 2025; Metz, 2022).

Brain-data covenant. This is an explicit organisational undertaking recognising neural data as a high-risk category. The covenant limits collection to strictly necessary purposes, prohibits opportunistic repurposing, imposes heightened security obligations, restricts transfer and monetisation, sets clear retention periods, and provides fair complaint and remedy mechanisms. In business settings, the covenant can be integrated into internal compliance regimes, ethics review boards, vendor contracts, and regulatory sandboxes. It operationalises the Ma'atian principle of reciprocity by converting organisational obligations into a binding stewardship commitment.

NUM councils. These oversight bodies combine legal, ethical, technical, labour, community, and sector-specific expertise to review high-impact uses of neurotechnology. Their functions include pre-deployment review, ongoing monitoring, complaint reception, impact assessment, and public accountability. NUM councils may be established at organisational, sectoral, or national level and are particularly important for high-risk domains such as mining, transport, health, and education (Eke, 2024; UNESCO, 2023). They provide the institutional anchor for communal responsibility, ensuring that the interests of affected communities are represented in governance decisions rather than assumed by technology developers or regulators alone.

Table 4. Institutional governance tools under NUM Ethics

Governance tool	Primary purpose	Minimum institutional effect	Addresses cognitive extractivism level
NUM design benchmark	Test necessity, proportionality, and dignity impact before deployment	Prevent casual or profit-driven neurotechnology adoption	Primary & secondary
Person-in-community consent	Make consent relationally intelligent but individually protective	Improve voluntariness, comprehension, and legitimacy	Primary
Brain-data covenant	Bind organisations to heightened neural data stewardship	Limit repurposing, monetisation, and downstream misuse	Primary, secondary & tertiary
NUM councils	Provide oversight, accountability, and multi-stakeholder scrutiny	Create durable review and redress pathways	All levels

9. Illustrative Policy Application

A brief policy illustration helps clarify the framework's practical use. Consider an employer in a mining context proposing fatigue-monitoring EEG headsets to reduce accident risk. A conventional compliance approach might ask only whether workers signed consent forms and whether some privacy notice was supplied. Under NUM Ethics, the inquiry would be substantially broader. The NUM design benchmark would ask whether EEG monitoring is strictly necessary relative to less intrusive safety measures such as rest schedules, supervisor assessment, or voluntary fatigue self-reporting. The person-in-community consent model would require clear

explanation of data use adapted to workers' literacy levels, opportunity for consultation with worker representatives or community intermediaries, and absolute preservation of individual refusal without disguised retaliation. The brain-data covenant would prohibit repurposing the data for productivity scoring, disciplinary surveillance, or commercial sale to insurers. A NUM council, composed of legal, labour, ethical, and community representatives, would review proportionality, retention limits, data security, and redress pathways before deployment approval is granted. This example illustrates how the framework shifts governance from formal permission to substantively ethical deployment, and from individual consent to communal legitimacy.

10. Discussion

The significance of NUM Ethics lies in its attempt to bridge philosophical reflection and institutional governance in a manner appropriate to African contexts. Rather than rejecting established principles such as mental privacy or cognitive liberty, the framework situates them within a broader account of relational personhood and communal responsibility drawn from African ethical traditions. In doing so, the article contributes to efforts to develop technology governance frameworks responsive to diverse moral traditions rather than derived exclusively from Western liberal models (Andanda & Düwell, 2024; Farahany, 2023; Metz, 2022).

The comparative analysis in Section 6 demonstrates that NUM Ethics is not merely a culturally particular variant of existing governance paradigms but a genuinely distinct fourth paradigm. Where rights-based models protect individual cognitive interests and market-based models default to commercial freedom, and state-centric models prioritise industrial coordination, the relational-communitarian paradigm addresses the governance of neuro-AI as a question of communal legitimacy, stewardship, and justice. This fourth paradigm speaks to a need identified across multiple governance domains: the need for frameworks that are not only formally correct but substantively resonant with the values and relational structures of the communities they govern.

The concept of cognitive extractivism sharpens this contribution by identifying the political economy problem that NUM Ethics is designed to resist. The concern is not simply privacy loss, but the emergence of institutional arrangements in which neural signals are appropriated, monetised, and operationalised under unequal conditions of power and benefit. This framing connects neuro-AI governance to broader debates about data colonialism, digital sovereignty, and epistemic justice in the Global South. At the same time, the expanded analysis of African data protection instruments in Section 4.5 corrects any impression that African governance is a blank slate: the continent has significant and growing data protection infrastructure, and NUM Ethics is proposed as a neuro-specific extension of that infrastructure rather than a replacement for it.

The framework is also relevant beyond Africa. As neurotechnology becomes globally distributed, governance models must engage with moral pluralism rather than assuming a single philosophical tradition is universally sufficient. The case for NUM Ethics therefore demonstrates how non-Western ethical traditions can contribute constructively to global governance design, especially in fields where emerging technologies create new forms of vulnerability and extraction. Policymakers developing domestic neurotechnology governance frameworks in other Global South contexts may find the relational-communitarian paradigm applicable, with appropriate adaptation.

At the same time, caution is necessary. NUM Ethics is not intended to replace existing legal regimes, nor does it claim to exhaust the diversity of African ethical thought. The framework is best understood as a heuristic and normative foundation for governance design. Its limitations are addressed in Section 12.

11. Policy Implications

For policymakers, the implications are immediate and concrete. Neural data should be expressly recognised in data protection or sector-specific legislation as a distinct sensitive category, not merely subsumed within existing biometric or health data definitions. The existing African data protection frameworks surveyed in Section 4.5 provide the legislative architecture within which this recognition can be introduced; amendment or supplementary regulation is the most practicable near-term pathway. Labour law and consumer protection law should address coercive neuro-monitoring and inferential profiling, particularly in high-risk sectors identified above. Public procurement frameworks should require NUM-style ethical due diligence before the adoption or importation of neurotechnology systems.

Regional African institutions have a particularly important role to play. The African Union, building on the Malabo Convention's foundational architecture, should consider developing a Model Protocol on Neurotechnology Data Governance that member states could incorporate into domestic law. SADC, ECOWAS, and EAC should explore sector-specific harmonisation frameworks, particularly for mining and transport contexts where neuro-surveillance is most immediately commercialising. Such regional instruments would enable African states to engage global neurotechnology market formation from a position of normative agency rather than reactive accommodation (Eke, 2024; UNESCO, 2025).

Viewed through the lens of cognitive extractivism, these reforms are not merely regulatory housekeeping. They are mechanisms for preventing the emergence of neural-data markets in which persons and communities become sources of cognitive value without corresponding dignity protection, bargaining power, or fair benefit-sharing. NUM Ethics provides policymakers with a vocabulary not only for governance design, but also for identifying and contesting inequitable political-economic arrangements.

A suggested phased implementation roadmap follows from this analysis. In the near term (2026–2027), African data protection authorities should issue interpretive guidance recognising neural data as a distinct category within existing sensitive data classifications. In the medium term (2027–2029), national legislatures should amend data protection acts to include explicit neural data provisions, cognitive liberty protections, and workplace neuro-surveillance restrictions. In the longer term (2029–2032), the African Union should develop and promote a Model Protocol on Neurotechnology Governance drawing on NUM Ethics principles, and sector-specific NUM councils should be piloted in high-priority domains. This roadmap is indicative rather than prescriptive; the pace and form of implementation will depend on national and regional political circumstances.

12. Limitations and Future Research

The framework proposed here remains primarily normative and requires empirical testing in sector-specific contexts. NUM Ethics makes governance proposals that are philosophically grounded and institutionally reasoned, but their practical effectiveness in specific African jurisdictions—different legal systems, institutional capacities, and socio-economic conditions—remains to be assessed. Future research should examine how NUM-style consent models operate in workplaces, health systems, education settings, and consumer markets across different African contexts.

The analysis relied principally on English-language sources, which may underrepresent Francophone, Lusophone, and Arabophone scholarship on both African ethics and neurotechnology governance. This is

especially significant for the Ma'at component, since relevant North African and Egyptological scholarship also appears outside English. The proposed framework should therefore be read as an analytically grounded but revisable synthesis. Future research should also examine whether the relational-communitarian paradigm can be usefully extended to other Global South governance contexts beyond Africa, and whether cognitive extractivism can be operationalised as a legal or regulatory category in domestic and international law. Comparative work across African jurisdictions on the legislative amendment pathways for neural data classification would also be valuable.

13. Conclusion

This article has argued that the rapid commercialisation of neurotechnology presents a significant governance challenge for African business and policy landscapes. Existing legal and ethical frameworks remain inadequate because they fail to recognise the exceptional status of neural data, the emerging realities of workplace and market-based neuro-surveillance, and the centrality of relational and communal values in many African contexts (Botes et al., 2025; Eke, 2024; Farahany, 2023). Through a structured review of the literature and a comparative analysis of global neurotechnology governance paradigms, the article has shown that African governance in this domain is not absent but is data-centric, institutionally uneven, and not yet neuro-specific—and that this precise gap calls for the framework developed here (Matshabane et al., 2024).

In response, the article has introduced Neuro-Ubuntu Ma'at (NUM) Ethics as a synthetic paradigm integrating Ubuntu, Ma'at, and broader indigenous African business ethics into a coherent framework for neuro-AI governance. It has also advanced the concept of cognitive extractivism at three analytical levels—primary, secondary, and tertiary—to describe the political economy of neural data extraction under conditions of weak regulatory oversight and asymmetrical technological power. It has positioned NUM Ethics as a fourth, relational-communitarian paradigm in a global typology of neurotechnology governance, demonstrating five specific comparative advantages over rights-based, market-based, and state-centric alternatives. And it has grounded the framework in existing African data protection developments, positioning it as an extension of rather than departure from current continental governance momentum.

The broader implication is that African-centred approaches to neuro-AI governance are not only normatively compelling but also institutionally viable. By drawing on indigenous ethical traditions alongside contemporary regulatory debates, policymakers and firms can develop governance architectures that protect cognitive liberty, strengthen accountability, and prevent exploitative forms of neural data extraction. As neurotechnology moves from experimental contexts into commercial and workplace environments across Africa and the Global South, the development of such frameworks will be essential to ensuring that technological innovation advances human dignity rather than undermining it.

Funding Statement

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of Conflicting Interests

The author declares no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Acknowledgements

The author declares that this manuscript was written independently. Grammarly was used solely for language editing, grammar checking, and improvements to clarity. No generative artificial intelligence system was used to produce the intellectual content, arguments, analysis, or findings of this paper. All conceptual framing, interpretation, and conclusions are the author's own.

References

- Andanda, P., & Düwell, M. (2024). Individualistic versus relational ethics – A contestable concept for (African) philosophy. *Ethical Theory and Moral Practice*. <https://doi.org/10.1007/s10677-024-10445-0>
- Beauchamp, T. L., & Childress, J. F. (2019). *Principles of biomedical ethics* (8th ed.). Oxford University Press.
- Botes, M., Labuschaigne, M., Casteleyn, C., Inkster, B., & Sheppard, B. (2025). Decoding the brain, respecting the person: A neuroethical inquiry into consent and cognitive liberty in South Africa. *Neuroethics*. Advance online publication. <https://doi.org/10.1007/s12152-025-09615-3>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Cheung, E. (2026, March 13). In a first for China, Neuracle's implantable brain-computer interface wins approval. *South China Morning Post*. <https://www.scmp.com/business/banking-finance/article/3346495/first-china-neuracles-implantable-brain-computer-interface-wins-approval>
- Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press.
- De Sousa Santos, B. (2018). *The end of the cognitive empire: The coming of age of epistemologies of the South*. Duke University Press.
- Eke, D. O. (2024). Ethics and governance of neurotechnology in Africa: Lessons from AI. *JMIR Neurotechnology*, 3, e56665. <https://doi.org/10.2196/56665>
- Ewuoso, C., & Hall, S. (2019). Core aspects of ubuntu: A systematic review. *South African Journal of Bioethics and Law*, 12(2), 93–103. <https://doi.org/10.7196/SAJBL.2019.v12i2.679>
- Farahany, N. A. (2023). *The battle for your brain: Defending the right to think freely in the age of neurotechnology*. St. Martin's Press.
- Gyekye, K. (1995). *An essay on African philosophical thought: The Akan conceptual scheme* (Rev. ed.). Temple University Press.
- Ienca, M., & Andorno, R. (2017). Towards new human rights in the age of neuroscience and neurotechnology. *Life Sciences, Society and Policy*, 13(1), Article 5. <https://doi.org/10.1186/s40504-017-0050-1>
- Karenga, M. (2004). *Maat, the moral ideal in ancient Egypt: A study in classical African ethics*. Routledge.
- Kenya Data Protection Act, No. 24 of 2019.
- Matshabane, O. P., et al. (2024). Advancing neuroethics in Africa. *South African Journal of Science*, 120(5/6), Article 18180. <https://doi.org/10.17159/sajs.2024/18180>
- Metz, T. (2022). *A relational moral theory: African ethics in and beyond the continent*. Oxford University Press.
- Molefe, M. (2019). *An African philosophy of personhood, morality, and politics*. Palgrave Macmillan.
- Musole, E. (2026). Governing neural data in the Global South: Toward a neuro-responsible business framework for Zambia. *International Journal of Advanced Business Studies*, 5(1), 47–61. <https://doi.org/10.59857/g3bvjx74>
- Ndlovu-Gatsheni, S. J. (2018). *Epistemic freedom in Africa: Deprovincialization and decolonization*. Routledge.
- Nigeria Data Protection Act, No. 32 of 2023.

- Obenga, T. (2004). African philosophy: The pharaonic period, 2780–330 BC. Per Ankh.
- OECD. (2019). Recommendation of the Council on Responsible Innovation in Neurotechnology. OECD. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0457>
- Pati, D., & Lorusso, L. N. (2018). How to write a systematic review of the literature. *HERD: Health Environments Research & Design Journal*, 11(1), 15–30. <https://doi.org/10.1177/1937586717747384>
- Protection of Personal Information Act 4 of 2013 (South Africa).
- Reuters. (2026, March 13). China approves market launch of brain-computer interface medical device in world first. Reuters. <https://www.reuters.com/business/healthcare-pharmaceuticals/china-approves-market-launch-brain-computer-interface-medical-device-world-first-2026-03-13/>
- Rwanda, Law No 058/2021 on the Protection of Personal Data and Privacy (Official Gazette of 25 October 2021).
- Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>
- UNESCO. (2023). An ethical framework for neurotechnology. UNESCO.
- UNESCO. (2025). AI and digital transformation in Africa. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000391406>
- Wiredu, K. (1996). Cultural universals and particulars: An African perspective. Indiana University Press.
- Xinhua. (2026, March 5). China identifies brain-computer interfaces as an industry of the future in government work report. Xinhua. http://en.ce.cn/Insight/202603/t20260303_2801486.shtml
- Yuste, R., Goering, S., Arcas, B. A. y, Bi, G., Carmena, J. M., Carter, A., Fins, J. J., Friesen, P., Giordano, J., Haggerty, J., Illes, J., Kellmeyer, P., Lee, W., Meyer, M. N., Pustilnik, A., & Wolpaw, J. (2017). Four ethical priorities for neurotechnologies and AI. *Nature*, 551(7679), 159–163. <https://doi.org/10.1038/551159a>
- Zambia Data Protection Act, No. 3 of 2021.
- Colorado Privacy Act, Colo. Rev. Stat. § 6-1-1301 et seq. (2021, as amended 2024) (neural data protection provisions).
- Illinois Biometric Information Privacy Act, 740 ILCS 14 (2008).