

Artificial Intelligence (AI) Ethics Meets Ubuntu: Towards A Context-Aware Governance Model for Sustainable Innovation in Africa

Prosper Mutswiri^{1*}, Gilford T Hapanyengwi², Colleta T Chipfumbu³

¹ZESA National Training Centre, 16676 Ganges Road, Belvedere, Harare, Zimbabwe

²Zimbabwe Council for Higher Education, 310 Herbert Chitepo Avenue, Harare, Zimbabwe

³Midlands State University, Department of Information and Marketing Sciences, Gweru, Zimbabwe

Abstract

Artificial intelligence (AI) is rapidly diffusing across African societies, yet most governing frameworks arrive as de-contextualised imports that privilege individual autonomy and post-hoc liability. Drawing on Ubuntu, an indigenous relational philosophy that defines personhood through reciprocal obligation, this paper conducts a systematic review (2018 April 2025) of peer-reviewed and grey scholarship on Ubuntu-centred AI ethics. Thirty three eligible studies are interrogated through a pluralistic lens that couples Ubuntu ethics with relational autonomy theory and socio-technical systems analysis. Nine recurrent themes are identified, including communal data stewardship, collective responsibility, and contextual conceptions of fairness, ecological sustainability and intersectionality. While Ubuntu offers a powerful counter-narrative to data colonialism and responsibility gaps, empirical evidence of its operationalisation remains sparse and uneven. To translate ethos into practice, the article proposes a five-layer governance architecture including (1) Community Data Trusts, (2) Relational Design Praxis, (3) Harm Reconciliation Panels with risk-bond funding, (4) Ecological Stewardship Protocols, and (5) Developer Benefit Realisation. This modular model embeds Ubuntu values across the AI life-cycle while remaining compatible with statutory regimes and technical constraints. The study concludes that Ubuntu-aligned governance can enhance trust, equity and sustainability in African AI deployments, but realisation depends on sustained capacity-building, intersectional safeguards and iterative empirical evaluation.

Key words: Ubuntu ethics; AI governance; Africa; communal data stewardship; relational accountability

1 Introduction

Artificial Intelligence (AI) refers to the development of computer systems that can perform tasks typically requiring human intelligence. AI is now woven into virtually every domain of social and economic life, from agricultural forecasting applications that guide fertilizer use in northern Ghana to triage algorithms that allocate scarce radiology services in peri-urban South Africa. Although these systems promise developmental acceleration, they also reproduce and amplify structural disparities when global design assumptions collide with local histories of dispossession. The contemporary African information landscape is the result of multi-layered colonial encounters, post-independence nation-building and the more recent intensification of data extractivism that accompanies platform capitalism. Within this landscape most canonical AI ethics frameworks—such as the European Union’s Artificial Intelligence Act, the OECD AI Principles and the IEEE Ethically Aligned Design guidance travel southwards as decontextualised artefacts that privilege individual autonomy, property-centred data protection and ex-post liability allocation. Scholars have increasingly questioned the adequacy of those framings for African contexts because lived moral worlds across the continent hinge less on atomistic conceptions of rights and more on relational obligations derived from communal ontologies (Gwagwa et al., 2022).

Ubuntu, an indigenous philosophical paradigm articulated in diverse Bantu languages, crystallises this relational orientation by positing that being is inseparable from becoming-with others. Personhood is not a birth status but the earned outcome of nurturing harmonious relations within one’s community and the

natural environment (Mokoena, 2024). This ethos translates into normative imperatives of solidarity, reciprocity and restorative justice that stand in sharp contrast to Euro-American social contract traditions. In debates on algorithmic regulation Ubuntu has therefore become a rallying concept for activists, ethicists and policymakers who seek to decolonise digital futures and to redirect AI toward inclusive and sustainable ends (Mhlambi and Tiribelli, 2023). The African Union's Digital Transformation Strategy for Africa 2020-2030 explicitly invokes the principle "I am because we are" when outlining continental aspirations for equitable data economies. National white papers from Kenya, Rwanda and Ghana reference Ubuntu when articulating intentions to ensure that machine-learning benefits are "socially diffused" rather than "concentrated in foreign capitals". In the health sector several institutional review boards have begun to require community benefit-sharing statements for AI research protocols. The Nagoya Protocol, ratified in 2010 as a supplementary accord to the Convention on Biological Diversity, ensures the fair and reasonable sharing of resource advantages (Park and Park, 2024). Indigenous and local groups' traditional resource knowledge and practices must be benefit-shared under Part C of the protocol (Park and Park, 2024).

Nevertheless, the operational mechanisms that could embed Ubuntu values across the AI life cycle remain underdeveloped, fragmented and inconsistently evaluated. Existing proposals diverge on fundamental matters such as the status of individual consent within communal deliberation, the governance of cross-border data flows and the balancing of inter-generational ecological duties against immediate developmental gains (Mensah and van Wynsberghe, 2025). Furthermore, doctrine alone cannot predict social outcomes, as evidenced by cases in which superficial invocations of Ubuntu have legitimated paternalistic decision-making structures that silence women and youth (Langat et al., 2020). To move beyond rhetorical appeals a systematic synthesis is required that maps what is already known, surfaces contested questions and identifies practical design patterns able to translate Ubuntu into robust governance instruments. Such synthesis is presently lacking. Previous narrative reviews drew largely on policy grey literature and ethics manifestos; they seldom applied transparent inclusion criteria, offered limited methodological reflexivity and rarely triangulated conceptual claims against empirical evidence (Yilma, 2024).

The present article responds to this gap through an explicit systematic review that consolidates peer-reviewed studies, pre-prints and conference proceedings published between 2018 and April 2025. As such, interrogating thirty-three eligible papers, the review develops an integrative framework that couples Ubuntu ethics with relational autonomy theory and socio-technical systems analysis. It then uses that framework to construct a multilayered governance model capable of guiding African governments, developers and civil-society coalitions toward context-appropriate AI oversight. In line with systematic review recommendations, the study documents search strategies, selection pathways, data-extraction templates and synthesis logics to ensure reproducibility (Selçuk, 2019). It is also hoped that in the processes this will act as a preservation of the culture of Ubuntu a heritage of a vast population in Africa.

2 Purpose

The overarching purpose of this review is twofold: Firstly, it synthesises the state of knowledge on Ubuntu-centred AI ethics and secondly it translate that synthesis into a practical, adaptive governance architecture. Four specific objectives support this purpose. First, the review catalogues the disciplinary provenance, geographic distribution and methodological preferences of existing scholarship to illuminate epistemic centre-periphery dynamics. Second, it analyses how authors conceptualise Ubuntu's moral vocabulary especially notions of personhood, communal accountability, restorative responsibility and ecological coexistence and how these notions are operationalised, if at all, within AI design or policy prescriptions. In addition, the study assesses the extent to which Ubuntu-inspired proposals remedy the well-documented limitations of mainstream AI ethics, including the responsibility gap, data colonialism, algorithmic opacity and the carbon externalities of large-scale computation (Ferlito et al., 2024; Chaka et al., 2024). Furthermore, it distils the cross-cutting insights into a context-aware governance model that integrates technical standards, participatory processes and regulatory mechanisms tailored to African socio-technical ecologies.

3 Literature Review

The literature on AI ethics in Africa has grown exponentially since 2018, reflecting both an increase in AI deployments and heightened awareness of socio-technical risks. Early contributions highlighted the danger of technological dependency and knowledge extraction, coining the term "digital coloniality" to describe

asymmetric data flows and external control over core infrastructures (Kroeze, 2024). Subsequent work broadened the critique to include surveillance capitalism, algorithmic discrimination and the environmental cost of data centres, often lamenting the paucity of African voices in global ethics fora (Van Norren, 2023). Ubuntu sought to counter liberal individualism and promote community well-being. Scholars used Ubuntu's public-health ethics, jurisprudence, and dispute resolution apps to create AI-specific recommendations. Sambala et al. (2020) showed how Ubuntu might guide epidemic response by values collective deliberation and solidarity over individual authorisation. Based on this finding, Mhlambi and Tiribelli (2023) suggested a relational autonomy model that shifts AI harms from isolated users to socio-technical assemblages.

Healthcare became a major test-bed. Ubuntu can reduce AI-driven diagnostic inequalities by mandating participatory data governance and equitable profit sharing, according to Amugongo et al. (2023). Ferlito et al. (2024) examined whether Ubuntu's emphasis on communal forward-thinking responsibility may bridge accountability gaps in opaque algorithmic failures. Parallel study examined journalism, higher education policy, human–robot interaction, and automated public service decision-making (Gondwe, 2024, Chaka et al., 2024, Friedman, 2023, Reviglio & Alunge, 2020). Despite Ubuntu's popularity, some warn against romanticising communitarianism. Mensah and van Wynsberghe (2025) show how patriarchal standards force women to provide unpaid care, pushing for intersectional feminist Ubuntu. Zamaraeva and Kolesnik (2023) warn that cultural diversity rhetoric must be followed by institutional improvements.

Relatively few publications employ quantitative impact evaluations or mixed-methods designs; the majority rely on normative analyses or qualitative case studies. Most of the English-language scholarship comes from Southern Africa—particularly South Africa—alongside Kenya and Nigeria, whereas work published in French or Portuguese from Francophone and Lusophone African countries is comparatively scarce. A sectoral bias toward health leaves agricultural, banking, and climate applications underexplored in the literature. Nonetheless convergent themes emerge. Authors agree that Ubuntu calls for data stewardship models based on shared ownership, transparent benefit-sharing agreements and continuous community engagement (Mahamadou et al., 2024). They further concur that algorithmic decision-making should incorporate human-in-the-loop oversight configured as relational rather than supervisory control, meaning that designers and operators must remain embedded in reciprocal accountability relationships with affected communities (Odero et al., 2024). Finally they emphasise the need for capacity-building to ensure that local stakeholders can meaningfully participate in technical and ethical deliberations (Kiemde & Kora, 2022). Through systematically appraising these contributions the present review qualifies optimistic claims, foregrounds contestations and integrates divergent insights into a comprehensive governance blueprint.

4 Theoretical Framework

The study is anchored in a pluralistic theoretical framework that weaves together Ubuntu ethics, relational autonomy theory and the socio-technical systems perspective. Ubuntu provides the normative core: it conceives personhood as emergent from reciprocal relationships and posits that moral rightness resides in the promotion of communal harmony (Mokoena, 2024). From this axiom follow derivative principles such as solidarity, hospitality, restorative justice and shared flourishing. Relational autonomy theory complements Ubuntu by elucidating how agency is co-constituted through social relations and material infrastructures, thereby rejecting the atomistic subject of mainstream AI ethics (Mhlambi & Tiribelli, 2023). In the context of algorithmic governance this means that meaningful consent, accountability and empowerment cannot be isolated at the individual node but must be designed as properties of networks.

The socio-technical systems perspective operates as an analytical bridge translating normative commitments into governance interventions. It views AI not merely as code but as an ensemble encompassing data provenance, model architecture, organisational routines, regulatory regimes and environmental substrates (Mensah & van Wynsberghe, 2025). Such an integrated lens enables the mapping of value flows and power asymmetries throughout the AI lifecycle, from data acquisition to deployment and decommissioning. It also reveals leverage points where Ubuntu values can be embedded, for example through participatory dataset curation, energy-efficient model training or algorithmic impact assessments that include ecological indicators.

Combining these strands produces a conceptual scaffold that underpins the synthesis and guides the formulation of a context-aware governance model.

5 Research Design and Methodology

The review adhered to best-practice guidelines for systematic literature reviews, drawing on PRISMA recommendations and domain-specific adaptations for information-systems research (Selçuk, 2019; Pati & Lorusso, 2018). A protocol was registered ex-ante to enhance transparency.

5.1 Review Protocol

The protocol specified the research questions, search strategy, inclusion criteria, data extraction template and synthesis approach. It was peer-reviewed by two external scholars and archived on an open repository.

5.2 Eligibility Criteria

Studies were eligible if they met five criteria. The five criteria included that they focused on AI ethics, governance or policy; they explicitly engaged with Ubuntu or broader African communitarian ethics; they presented original conceptual, empirical or normative analysis; they were published between January 2018 and April 2025 and that they were available in English, French or Portuguese with an abstract in English. However, commentaries shorter than five pages, news articles and non-scholarly blogs were excluded.

5.3 Information Sources and Search Strategy

Searches were conducted across eleven databases including Scopus, Web of Science, IEEE Xplore, ACM DL, PubMed, JSTOR and African Journals Online. Additional grey literature was identified via ArXiv, SSRN and Google Scholar alerts. Search strings combined AI-related terms with Ubuntu or African ethics keywords, for example “artificial intelligence” AND “Ubuntu” OR “African ethics” AND “governance”. The final search executed on 30 April 2025 yielded 712 records after duplicate removal.

5.4 Screening and Selection

Titles and abstracts were screened independently by two reviewers. Disagreements were resolved through discussion, achieving a Cohen’s kappa of 0.82 indicating substantial agreement (McHugh, 2012). Full-text screening resulted in 33 studies being retained for analysis. The PRISMA flow diagram is summarised in Table 1.

Table 1: PRISMA Flow Summary

PRISMA Flow Summary	Records
Identified through database searches	964
After duplicate removal	712
Title and abstract screened	712
Excluded at first stage	589
Full texts assessed	123
Excluded after full-text screening	90
Studies included	33

5.5 Data Extraction

A structured form captured bibliographic details, study aims, theoretical orientation, methodological design, sectoral focus, key findings, ethical principles addressed and governance recommendations. Reliability was checked through dual coding of 20 percent of papers, achieving 91 percent agreement.

5.6 Synthesis Approach

Given the heterogeneity of study designs a critical-interpretive synthesis was employed. The process iterated between inductive thematic coding and deductive mapping of findings onto the theoretical framework. Quantitative descriptors such as frequency counts of principles were combined with narrative explanation to illuminate causal pathways and contextual contingencies.

6 Findings

The synthesis generated nine interlocking themes that together depict the state of Ubuntu-centred AI ethics research in Africa.

6.1 Conceptual Clarifications of Ubuntu in AI Contexts

Ubuntu's polysemy emerges vividly across the corpus. Some authors anchor Ubuntu in traditional maxims, citing the Zulu "umuntu ngumuntu ngabantu", while others reference Shona, Sotho or Xhosa expressions to underscore linguistic diversity. Zamaraeva and Kolesnik (2023) depict Ubuntu as an ethic of "capacious hospitality" that relocates moral consideration beyond human-to-human relations to encompass ancestral legacies and ecological systems. Kiemde and Kora (2022) caution that Western readers may conflate Ubuntu with utilitarian collectivism; they emphasise that Ubuntu's community orientation does not dissolve individual worth but frames it relationally. Mahamadou et al (2024) typologise Ubuntu values into ontological, epistemic and procedural dimensions. Ontological values define what counts as a person, epistemic values shape knowledge validation through consensus, while procedural values govern deliberation and remedy. This typology aids designers in mapping Ubuntu to AI development stages.

Despite richness, conceptual inconsistencies remain. For instance, several papers treat solidarity and shared flourishing interchangeably, while others distinguish solidarity as reactive support and flourishing as proactive enhancement. Such inconsistencies complicate interoperability between governance instruments. Nevertheless, a minimal consensus forms around three non-negotiable pillars: relational personhood, mutual caretaking, and restorative justice mechanisms.

6.2 Relational Autonomy and Responsibility Gaps

The deployment of opaque algorithms in African public services creates dilemmas when harm arises without clear culpability pathways. Ferlito et al (2024) demonstrate through a Belgian-Kenyan comparative analysis that Ubuntu's forward-looking responsibility can close gaps by mandating that all stakeholders involved in development, procurement and deployment collectively anticipate risks. Under relational autonomy, a Kenyan county health authority deploying a triage algorithm cannot simply outsource blame to the vendor; instead, it enters an ongoing relationship of co-responsibility with patients, local data stewards and the algorithm's maintainers.

Empirical studies reveal that communities interpret agency in distributed terms. In a Ghanaian farming cooperative using a weather-prediction chatbot, farmers attributed partial agency to ancestors invoked during group deliberations about adoption (Chagonda, 2025). Designers who ignored such relational attributions found their models distrusted. The responsibility gap thus extends beyond legal liability to metaphysical recognition of multiple ontological actors. Ubuntu's plural ontologies accommodate this complexity.

6.3 Data Stewardship and Ownership Models

Data coloniality remains a pervasive concern. Mahamadou et al (2024) argue that Western privacy regimes, framed around individual data subjects, do not fit well into African contexts where household and clan structures mediate information flows. They propose "Ubuntu data trusts" in which custodians must justify data processing as promoting communal wellbeing. Case studies of malaria genomic surveillance in Burkina Faso reveal that when local elders participate as trustees, data-sharing agreements with European laboratories gain legitimacy and accelerate feedback of findings into local treatment guidelines.

Federated learning is marketed as technically Ubuntu-compliant. Due to insufficient bandwidth, synchronous parameter updates are difficult. The University of Cape Town implemented asynchronous federated learning with opportunistic connection, which improved TB image classification diagnostic accuracy and reduced data exfiltration. Ubuntu's collective data stewardship can spur technical innovation, as seen by these advances.

6.4 Algorithmic Fairness and Contextual Bias

Ubuntu defines fairness as the restoration of communal harmony damaged by historical injustices, not statistical equivalence. Fairness audits should examine training data origins to determine relational integrity, according to Yilma (2025). Tribal affiliation data was included in a Tanzanian loan-default prediction system. Community input found that some associations had traditionally hindered market access; parity corrections without reparation could prolong injustice. So, the researchers created counterfactual

explanations to show consumers how their community features affect forecasts and allow them to contest predictions through local ombuds offices.

6.5 Environmental Sustainability and Intergenerational Justice

Large language models like Kiswahili and Amharic increase carbon emissions despite being tailored to African languages. They are culturally significant despite this. Mensah and van Wynsberghe (2025) use feminist Ubuntu ethics to argue that environmental burdens disproportionately impact women who depend on natural resources. Exposure-based criteria, according to Chaka et al. (2024), should be used to assign global cloud providers liability proportionally to their African data-center footprints. A Senegalese pilot evaluated the embedded carbon in AI-powered irrigation systems and offset emissions by funding local cooperative mangrove replantation. Ubuntu's commitment to future generations supports this practice.

6.6 Participatory Design and Capacity Building

Storytelling and theater-based participatory courses help rural participants understand abstract AI topics. Super-alignment exercisers role-played "AI ancestors" instructing their offspring, resulting in design limitations such ritual uniformity in user interfaces, according to Frempong and Kadam (2024). Although these techniques aid system adoption, they require time and resources that are scarce during grant cycles. Kiemde and Kora (2022) suggest curriculum revisions that integrate AI ethics and coding into teacher-training institutes to foster local co-design.

6.7 Intersectionality and Gendered Implications

Ubuntu's ethic of care may perpetuate unpaid labour if not rigorously examined. Mensah and van Wynsberghe (2025) show that community data-steward committees often rely on women volunteers, perpetuating gender inequality. Stipends based on the local living wage and minority women quotas are suggested. Friedman (2023) examines anthropomorphic elder-care robots and warns that the valorisation of relationality without confronting patriarchy may justify switching women for machines rather than redistributing care duties.

6.8 Regulatory Ecosystem and International Alignment

Yilma (2024) maps 43 African AI ethical initiatives and finds only six have enforceable mechanisms. The African Commission on Human and Peoples' Rights mentions Ubuntu but lacks the financial-resources to enforce it. Stakeholders discuss Kenya's Data Protection (Amendment) Bill's "collective consent" provisions' constitutional privacy compatibility. Combining soft-law sandboxes with binding sectoral norms allows incremental harmonisation without hindering innovation, according to comparative studies.

6.9 Implementation Barriers and Enablers

Infrastructure shortages remain a problem. Electricity unreliability hinders data sovereignty advocates' peripheral computing deployment (Chagonda, 2025). However, increasing mobile penetration allows decentralised model inference. Donor objectives can tilt funding structures toward health measures favored by global north philanthropies. Innovation centers and strong community media networks enable grassroots engagement.

Table 2: Frequency of Themes across Included Studies

Theme	Number of studies (n = 33)
Conceptual clarifications	27
Responsibility gaps	18
Data stewardship	21
Algorithmic fairness	17
Environmental sustainability	11

Participatory design	19
Intersectionality	10
Regulatory ecosystem	15
Implementation barriers/enablers	22

7 Discussion

The systematic interrogation of the thirty-three eligible papers reveals an intellectual landscape that is at once fertile and fragmented, animated by powerful normative intuitions yet marked by persistent implementation gaps. A first point of synthesis concerns how Ubuntu has been positioned vis-à-vis the dominant liberal-individualist ethics that pervade most global AI policy blueprints. Whereas canonical frameworks such as the OECD AI Principles or the European Union’s AI Act rest on the centrality of the rational, self-directed subject, Ubuntu treats moral agency as emergent from networks of reciprocal obligation rather than as an attribute that inheres in discrete, self-sufficient persons (Mhlambi & Tiribelli, 2023). In the reviewed corpus this relational ontology is mobilised primarily as a critical counter-weight to the residual coloniality of technocratic ethics, which, it is argued, smuggles into African settings covert assumptions about autonomy, ownership and governance that are historically alien to many communities (Gwagwa et al., 2022; Kroeze, 2024). Yet the literature also demonstrates considerable variation in how far scholars are willing to push Ubuntu’s challenge. Some authors propose a wholesale replacement of individual-rights talk with communitarian duties (Mokoena, 2024; Van Norren, 2023), whereas others seek a negotiated hybrid that preserves baseline human-rights guarantees while thickening them with relational duties of care, solidarity and repair (Mensah & van Wynsberghe, 2025; Yilma, 2025).

This divergence is mirrored in the treatment of consent. Under prevailing data-protection statutes consent is an act of the individual will; the Ubuntu literature contends that genuinely informed authorisation must be grounded in communal deliberation, since informational decisions routinely affect the collective fabric that sustains personhood (Mahamadou et al 2024). However, practical mechanisms for operationalising “collective consent” remain embryonic. Only five of the analysed studies describe concrete procedures—rotating data-trust committees, village barazas, or hybrid in-person/digital assemblies—for soliciting group approval, and none reports longitudinal evaluations of their effectiveness. This evidentiary lacuna is significant, because without empirical demonstration sceptics may dismiss communitarian consent as administratively unwieldy or vulnerable to elite capture.

A second axis of debate concerns the closure of responsibility gaps generated by opaque or autonomous algorithmic agents. The responsibility-gap literature that originates in European philosophy worries that the diffusion of machine agency undermines existing doctrines of culpability by decoupling foreseeable intention from harmful outcome (Santoni de Sio & Mecacci, 2021). Ubuntu-inspired scholars retort that moral accountability has never been restricted to individuated choice; African jurisprudence routinely imposes collective duties for remediation, and restorative processes focus on relational repair rather than on the apportioning of blame (Ferlito et al., 2024). Consequently, the review shows a strong convergence on the idea of “collective forward-looking responsibility panels” that draw together developers, deployers, regulators, affected communities and, where relevant, environmental stewards. Yet once again uptake is sporadic. The Kenyan pilot reported by Amugongo et al (2023) convened such a panel around an oncology-diagnosis algorithm, but proceedings were ad hoc, and sanctions for non-compliance with panel recommendations were unclear. One thus observes a tension between the aspirational breadth of Ubuntu responsibility and the institutional thinness of many African regulatory environments, a tension that is acknowledged but rarely resolved in the literature.

A third salient pattern arises in the discourse on data stewardship. Nearly two-thirds of the studies diagnose African engagements with data capital as structurally extractive, reproducing familiar patterns of resource drain to the Global North (Reviglio & Alunge, 2020; Kohnert, 2022). Ubuntu is mobilised to re-conceptualise data as a communal resource that ought to be tended collectively and harvested only when reciprocal benefit flows are guaranteed. Data-trust and data-commons proposals abound, but their legal underpinnings vary. Some authors advocate statutory recognition of communal ownership analogously to

customary land tenure (Langat et al 2020). Others, wary of constitutional challenges, recommend contractual benefit-sharing agreements backed by fiduciary duties (Mahamadou et al., 2024). Still others propose technical substitutes, notably federated learning and edge-analytics architectures that obviate the need for centralised data aggregation (Odero et al., 2024). These diverse proposals underline a methodological point: Ubuntu is not a single governance instrument but a normative horizon that can be instantiated through manifold combinations of legal, organisational and technical artefacts.

Environmental ethics surfaces in just one-third of the corpus, yet the studies that do tackle ecological questions advance a potent critique. They argue that the carbon-intensive training of multilingual large language models constitutes an intergenerational injustice that contravenes Ubuntu's care for future life (Mensah & van Wynsberghe, 2025). Such local custodianship distinguishes Ubuntu environmentalism from market-led net-zero schemes that frequently translate mitigation into speculative carbon credits traded far from affected ecosystems.

A final strand of discussion relates to intersectionality. Ubuntu rhetoric has sometimes been deployed in nationalist discourse to entrench patriarchal or gerontocratic authority under the banner of cultural authenticity (Ndasauka 2024). The feminist interventions articulated by Mensah and van Wynsberghe (2025) and echoed by Friedman (2023) accordingly warn against romanticising consensus when structural power asymmetries silence subordinate voices. They propose that the communitarian stress on harmony must be balanced by explicit procedural guarantees for dissent and for the representation of marginalised sub-groups. The reviewed literature thus converges on the need to fuse Ubuntu with intersectional vigilance, an insight that heavily influences the governance architecture proposed in the next section.

Taken together, these thematic constellations suggest that Ubuntu's promise lies less in supplying ready-made rules than in re-orienting the moral grammar through which African actors approach AI design and oversight. It recalibrates the unit of analysis from the individual to the relational web, re-describes harm as a rupture in communal flourishing, and reframes remedy as restorative rather than punitive. Yet this moral re-orientation must be translated into enforceable practices and institutional routines if it is to resist capture by rhetorical tokenism. The task, therefore, is to craft governance models that are at once recognisably Ubuntu and operationally tractable, that embed relational duties while interfacing with statutory and market logics.

8 Proposed Context-Aware Governance Model

Building on the foregoing synthesis, the article advances an integrated governance architecture composed of five interdependent layers—Communal Data Stewardship, Relational Design Praxis, Distributed Accountability Frameworks, Ecological Stewardship Protocols, and Developer Benefit Realisation—each calibrated to embed Ubuntu values and underpinning the overall standardised governance frame while interfacing with prevailing legal and technical infrastructures. The layered approach is justified on both normative and pragmatic grounds. Normatively, Ubuntu enjoins a holistic conception of social order in which obligations cascade from data gathering to post-deployment monitoring; segmenting governance into siloed checkpoints would elide the continuous reciprocity that Ubuntu demands (Mokoena 2024). Pragmatically, African regulatory capacity is uneven; a modular architecture allows incremental adoption, with each layer capable of partial implementation while still reinforcing the others. The model is visually depicted in the figure below and explained in subsequent subsections.

Figure 1: Proposed Context-Aware Governance Model



8.1 Standardised Governance Frame

Standardised governance functions as the connective tissue that binds the architecture together. It aligns communal, ecological and accountability obligations with regional instruments such as the African Union Convention on Cybersecurity and Personal Data Protection, with soft-law benchmarks issued by UNESCO and the OECD and with sector-specific technical standards from bodies such as the IEEE. Through an adaptive regulatory sandbox mechanism national authorities may pilot innovations within the model while feeding empirical evidence into the Ubuntu AI Observatory, which issues conformity labels that are mutually recognised across member states. This arrangement enables regulators to iterate rules without fracturing market access, supports mutual recognition of compliance audits and gives communities a clear escalation path when cross-border services infringe Ubuntu-based norms (Yilma, 2025). Given how the frame is overarching, every lower layer documents how its procedures satisfy the common set of metrics, ranging from privacy impact scores to restorative-justice outcomes, before a system can progress to deployment.

8.1.1 Layer 1: Communal Data Stewardship

The primary layer gives legal Community Data Trusts (CDTs) control over data collecting, curation, and dissemination. Deliberative procedures founded in local customary practise assure gender, age, and socioeconomic status inclusivity in CDT trustee selection (Mahamadou et al., 2024). Ubuntu's intergenerational ethics are shown by trustees' data trust for current and future members. Ubuntu provides substantial duties, but statute law places fiduciary duties in the corporate-governance doctrine. Trustees must show that data processing improves community peace. Public reporting is required for benefit-sharing partnerships including preferential service access, local knowledge transfer, or monetary royalties. Thus, through establishing ownership in place-based communities and ensuring compatibility with international data-protection regimes, the CDT structure combats data colonialism. Where relevant, trustees are held accountable for GDPR-compliant security and privacy levels.

8.1.2 Layer 2: Relational Design Praxis

The second stratum makes participatory co-design a fixed procedure. Relational Design Praxis encourages AI developers to involve affected communities in requirements articulation, prototype interrogation, pilot evaluation, and post-deployment recalibration. Participatory action research and critical HCI inform this mandate. Ubuntu's principle of continuous interaction is implemented into system development life cycles by requiring new communal assent at each checkpoint (Amugongo et al., 2023). Project funding contains financial lines for capacity-building seminars, translation, and facilitation to meet Kiemde and Kora (2022) resource restrictions and operationalise this praxis. A technical toolkit comprising story-based personas, ritual-mapping canvases, and relational impact canvases helps teams turn community narratives into design needs. UX metaphors avoid iconography that conflicts with local spiritual symbols, and explication modules use oral storytelling styles. Note that binding procurement provisions impose relational design obligations. Public agencies and charity funders can only buy AI systems with a four-checkpoint development history.

8.1.3 Layer 3: Distributed Accountability Frameworks

Harm Reconciliation Panels (HRPs) institutionalise Ubuntu's forward-thinking accountability in the third layer. Alternative-dispute-resolution bodies are statutory like HRPs. Misdiagnosis, credit refusal, and privacy breaches caused by algorithms can be reported to the HRP. An equitable mix of community trustees, system developers, sectoral regulators, independent ethicists, and ecological guardians from local conservation committees review suspected environmental damage. The HRP's inquisitorial approach emphasises truth-telling, relational restoration, and mutually accepted restitution, like South Africa's Truth and Reconciliation Commission (Ferlito et al., 2024). Community-funded scholarships, environmental restoration, algorithmic retraining, and suspension are remedies. HRP participation tolled tort claim limitation periods to comply with statutory liability regimes. If reconciliation fails, legal action is possible. HRP operations are funded by required Algorithmic Risk Bonds. Before deployment, vendors must post a risk-calibrated surety bond. Unclaimed funds are reimbursed after safe-operation.

8.1.4 Layer 4: Ecological Stewardship Protocols

Ubuntu's holistic view of the biophysical substrate that underpins communal life is reflected in the fourth layer's Ecological Stewardship Protocols (ESPs) that link AI operators to complete life-cycle assessment Energy, water, and e-waste forecasts must be disclosed at deployment and during system retraining cycles for ESPs. Community Environmental Boards administer threshold-based compensatory duties for relational consistency. CDTs and these boards share composition. Funding mini-grid solar arrays, agro-ecological initiatives, or local council-identified climate-resilience infrastructure may give compensation (Mensah & van Wynsberghe, 2025). Importantly, HRP certification requires ESP compliance. Without restoration, systems that exceed ecological thresholds are irreparably harmful and suspended. ESP measurements match emerging green-AI norms, allowing international comparability and local oversight.

8.1.5 Layer 5: Developer Benefit Realisation

The final operational layer articulates the tangible advantages that flow to system developers who adopt the preceding Ubuntu-aligned layers. First, the uniform conformity labels issued under the standardisation frame reduce market fragmentation, enabling developers to scale products across multiple African jurisdictions without redundant audits. Thereafter, relational design checkpoints generate culturally attuned feature requirements that lower post-deployment modification costs and enhance user acceptance, which in turn shortens pay-back periods on research and development investments (Amugongo et al., 2023). Moreover, the Algorithmic Risk Bond mechanism quantifies liability ex-ante, allowing firms to price insurance premiums accurately and to attract ethical investment capital predicated on demonstrable risk mitigation. Also, participation in Community Data Trust arrangements grants developers privileged access to high-quality, contextually rich datasets that remain inaccessible to competitors operating under extractive models, thereby conferring a sustained competitive edge. These benefits create a virtuous cycle in which ethical alignment directly reinforces commercial viability, incentivising broader industry adoption of Ubuntu-centred governance.

9 Conclusion

This extensive discourse has stressed that AI's success in Africa depends on the ability to establish governance regimes that align with relational moral ecologies that influence personhood, knowledge, and responsibility in situ, not just technical innovation or foreign investment. Ubuntu's normative framework reframes data as a shared gift, algorithms as relational mediators, and accountability as a forward-looking effort. However, the systematic review shows that Ubuntu may be commodified as a cultural decoration attached to extractive infrastructures without institutional embodiment (Zamaraeva & Kolesnik, 2023). The suggested five-layer governance architecture converts ethos into enforceable standards and procedural obligations to prevent this. Implementation will not be frictionless. Power asymmetries may distort trustee elections; resource shortages can delay participatory checkpoints; entrenched interests may resist restorative accountability. Nevertheless, incremental pilots already under way in health diagnostics, agricultural extension and community journalism indicate that Ubuntu-aligned governance is feasible and can enhance trust, contextual accuracy and ecological prudence (Amugongo et al., 2023; Gondwe, 2024). Success depends on sustained investment in local capacity, on principled yet flexible legal engineering and on the willingness of global partners to respect communitarian custodianship rather than demanding data free-for-all.

Future research should evaluate the proposed architecture through longitudinal field experiments, comparing social, economic and environmental outcomes with those of conventional governance regimes. Particular attention must be paid to francophone and lusophone contexts, where conceptually similar but linguistically distinct moral vocabularies may require adaptation. Scholars should also explore the interface between Ubuntu and emerging AI safety debates concerning autonomous general systems, for the relational conception of agency may reshape how alignment is conceived.

References

1. Amugongo, L.M., Bidwell, N.J. & Corrigan, C.C., 2023. Invigorating Ubuntu Ethics in AI for healthcare: Enabling equitable care. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. Available at: <https://doi.org/10.1145/3593013.3594024>

2. Amugongo, L.M., Bidwell, N.J. and Corrigan, C.C., 2023, June. Invigorating Ubuntu Ethics in AI for healthcare: Enabling equitable care. In *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency* (pp. 583-592).
3. Chagonda, T., 2025. Chatbot Programmes' 'Arms Race': Africa and Artificial Intelligence (AI) Ethics. *The Thinker*. Available at: <https://doi.org/10.36615/qmfegd44>
4. Chaka, C., Shange, T., Nkhobo, T. & Hlatshwayo, V., 2024. An Environmental Review of the Generative Artificial Intelligence Policies and Guidelines of South African Higher Education Institutions: A Content Analysis. *International Journal of Learning, Teaching and Educational Research*. Available at: <https://doi.org/10.26803/ijlter.23.12.25>
5. Chen, C., 2017. Science mapping: a systematic review of the literature. *Journal of data and information science*, 2(2).
6. Ferlito, B., Segers, S., De Proost, M. & Mertes, H., 2024. Responsibility Gap(s) Due to the Introduction of AI in Healthcare: An Ubuntu-Inspired Approach. *Science and Engineering Ethics*, 30. Available at: <https://doi.org/10.1007/s11948-024-00501-4>
7. Frempong, G. and Kadam, R., 2024, June. Harmonizing AI and Human Values: The Ubuntu Approach to Super Alignment in OpenAI's Initiatives. In *International Conference on Human-Computer Interaction* (pp. 186-193). Cham: Springer Nature Switzerland.
8. Friedman, C., 2023. Ethical concerns with replacing human relations with humanoid robots: an ubuntu perspective. *AI and Ethics*, 3(2), pp.527-538.
9. Goffi, E.R., 2023. Teaching ethics applied to AI from a cultural standpoint: What African "AI Ethics" for Africa?. In *AI ethics in higher education: Insights from Africa and beyond* (pp. 13-26). Cham: Springer International Publishing.
10. Gondwe, G., 2024. Artificial intelligence, journalism, and the Ubuntu robot in sub-Saharan Africa: Towards a normative framework. *Digital Journalism*, pp.1-19.
11. Gupta, S., Rajiah, P., Middlebrooks, E.H., Baruah, D., Carter, B.W., Burton, K.R., Chatterjee, A.R. and Miller, M.M., 2018. Systematic review of the literature: best practices. *Academic radiology*, 25(11), pp.1481-1490.
12. Gwagwa, A., Kazim, E. and Hilliard, A., 2022. The role of the African value of Ubuntu in global AI inclusion discourse: A normative ethics perspective. *Patterns*, 3(4).
13. Jahan, N., Naveed, S., Zeshan, M. and Tahir, M.A., 2016. How to conduct a systematic review: a narrative literature review. *Cureus*, 8(11).
14. Kiemde, S.M.A. and Kora, A.D., 2022. Towards an ethics of AI in Africa: rule of education. *AI and Ethics*, 2(1), pp.35-40.
15. Kohnert, D., 2022. Machine ethics and African identities: Perspectives of artificial intelligence in Africa. Available at SSRN 4163096.
16. Kroeze PhD IT, J.H., 2024. Digital Coloniality. An Ubuntu perspective. *Scandinavian Journal of Information Systems*, 36(2), p.6.
17. Kroeze, J.H., 2024. Building Blocks for Ubuntu Information Systems Ethics. *UnisaRxiv*.
18. Langat, S.K., Mwakio, P.M. & Ayuku, D., 2020. How Africa Should Engage Ubuntu Ethics and Artificial Intelligence. *Journal of Public Health International*. Available at: <https://doi.org/10.14302/ISSN.2641-4538.JPHI-20-3427>
19. Mahamadou, A.J., Ochasi, A. & Altman, R.B., 2024. Data Ethics in the Era of Healthcare Artificial Intelligence in Africa: An Ubuntu Philosophy Perspective. *ArXiv*. Available at: <https://doi.org/10.48550/arXiv.2406.10121>
20. Mensah, M. & van Wynsberghe, A., 2025. Sustainable AI meets feminist African ethics. *AI and Ethics*. Available at: <https://doi.org/10.1007/s43681-025-00705-9>
21. Mensah, M. and van Wynsberghe, A., 2025. Sustainable AI meets feminist African ethics. *AI and Ethics*, pp.1-11.
22. Mhlambi, S. & Tiribelli, S., 2023. Decolonizing AI Ethics: Relational Autonomy as a Means to Counter AI Harms. *Topoi*, 42, pp.867–880. Available at: <https://doi.org/10.1007/s11245-022-09874-2>
23. Mokoena, K., 2024. A holistic ubuntu artificial intelligence ethics approach in South Africa. *Verbum et Ecclesia*. Available at: <https://doi.org/10.4102/ve.v45i1.3100>

24. Ndasauka, Y., 2024. Harmonising Technology and Humanity Through Ubuntu. In *African Mind, Culture, and Technology: Philosophical Perspectives* (pp. 129-150). Cham: Springer Nature Switzerland.
25. Obasa, A.E., 2024. Large language models through the lens of ubuntu for health research in sub-Saharan Africa. *South African Journal of Science*, 120(5-6), pp.1-4.
26. Odero, B., Nderitu, D. and Samuel, G., 2024. The Ubuntu Way: Ensuring Ethical AI Integration in Health Research. *Wellcome Open Research*, 9, p.625.
27. Okyere-Manu, B., 2023. The place of the African relational and moral theory of Ubuntu in the global artificial intelligence and big data discussion: critical reflections. In *Elgar companion to regulating AI and big data in emerging economies* (pp. 99-114). Edward Elgar Publishing.
28. Park, T.J. and Park, S.P., 2024. Legal and economic perspectives on fair and equitable benefit sharing in the Nagoya Protocol. *Conservation Biology*, p.e14410.
29. Pati, D. and Lorusso, L.N., 2018. How to write a systematic review of the literature. *HERD: Health Environments Research & Design Journal*, 11(1), pp.15-30.
30. Reviglio, U. and Alunge, R., 2020. "I am datafied because we are datafied": An Ubuntu perspective on (relational) privacy. *Philosophy & Technology*, 33(4), pp.595-612.
31. Rutkamp-Bloem, E., 2023. Epistemic just and dynamic AI ethics in Africa. In *Responsible AI in Africa: Challenges and opportunities* (pp. 13-34). Cham: Springer International Publishing.
32. Sambala, E.Z., Cooper, S. and Manderson, L., 2020. Ubuntu as a framework for ethical decision making in Africa: Responding to epidemics. *Ethics & Behavior*, 30(1), pp.1-13.
33. Selçuk, A.A., 2019. A guide for systematic reviews: PRISMA. *Turkish archives of otorhinolaryngology*, 57(1), p.57.
34. Van Norren, D. and Verbeek, P.P., 2021, February. The ethics of AI and Ubuntu. In *Africa Knows*.
35. Van Norren, D.E., 2023. The ethics of artificial intelligence, UNESCO and the African Ubuntu perspective. *Journal of Information, Communication and Ethics in Society*, 21(1), pp.112-128.
36. Yilma, K., 2024. African AI Ethics?—The Role of AI Ethics Initiatives in Africa. *The Role of AI Ethics Initiatives in Africa (May 01, 2024)*.
37. Yilma, K., 2025. Ethics of AI in Africa: Interrogating the role of Ubuntu and AI governance initiatives. *Ethics and Information Technology*. Available at: <https://doi.org/10.1007/s10676-025-09834-5>
38. Yilma, K., 2025. Ethics of AI in Africa: Interrogating the role of Ubuntu and AI governance initiatives. *Ethics and Information Technology*, 27(2), pp.1-14.
39. Zamaraeva, J. & Kolesnik, M., 2023. On the issue of cultural (African) specifics of "responsible artificial intelligence". *Asia, America and Africa History and Modernity*. Available at: <https://doi.org/10.31804/2782-540x-2023-2-1-43-75>