Sober Search

A Journal of Philosophy

Published by the Department of Philosophy, University of Abuja

Volume 1 (2025), Article 5, https://doi.org/10.70118/SSJP0005

Technology and Moral Responsibility: An African Critique of Algorithmic Ethics

Emmanuel C. Ilo
Department of Philosophy, University of Ibadan

Abstract

The proliferation of artificial intelligence (AI) systems across diverse sectors globally has provoked critical ethical inquiries concerning fairness, responsibility, and justice. Predominantly shaped by Western epistemological frameworks, prevailing algorithmic ethics often marginalise non-Western traditions and exclude African voices from global AI governance. This paper critiques the Eurocentric bias inherent in mainstream AI ethics by advancing an African philosophical response rooted in Ubuntu—an indigenous moral framework that emphasises communal personhood, relational autonomy, and shared responsibility. Through a conceptual-normative and hermeneutic methodology, the study synthesises philosophical literature, policy documents, and recent case studies to examine how Ubuntu can inform the ethical design, deployment, and regulation of AI technologies in African contexts. It further employs the Discovery-Translation-Verification (DTV) framework to operationalise Ubuntu principles in technical and governance settings. Findings reveal that Ubuntu offers a robust normative alternative capable of addressing testimonial and hermeneutical injustices, while reshaping algorithmic accountability through participatory and restorative models. The paper concludes by proposing actionable recommendations for embedding African ethical values into AI policy, regulation, education, and design processes. It argues that far from being a regional ethic, Ubuntu offers a globally relevant moral paradigm for constructing inclusive, dignified, and context-sensitive algorithmic systems.

Keywords: Ubuntu Philosophy, Algorithmic Ethics, Epistemic Injustice, African AI Governance

Introduction

Artificial Intelligence (AI) has emerged as one of the most transformative technological innovations of the twenty-first century, permeating a wide array of societal domains, including governance, healthcare, finance, education, agriculture, and security. Across the globe, AI systems are increasingly leveraged to optimise decisions, streamline bureaucratic operations, predict behaviours, and enhance economic efficiencies. However, this rapid integration of AI technologies has not occurred in a moral vacuum. The ethical, social, and political

ramifications of AI have triggered significant scholarly concern, particularly regarding questions of fairness, transparency, accountability, and justice. At the heart of these debates lies the recognition that algorithms, far from being neutral tools, are embedded with the values and biases of their creators, and thus reproduce the structural inequities and epistemic asymmetries of the societies from which they emerge. In the African context, this concern is especially urgent. Many AI systems deployed on the continent are designed and owned by Western corporations, often without meaningful engagement with local knowledge systems, cultural values, or social realities. Birhane (2020) aptly refers to this phenomenon as "algorithmic colonisation," capturing the continuities between digital technologies and historical patterns of imperial domination. For instance, tech multinationals frequently roll out AI-driven services across African countries – from biometric ID systems to natural language processing tools – without adapting their models to African sociolinguistic diversity or ethical norms. Such practices not only perpetuate techno-dependency but also risk epistemic erasure, whereby indigenous ways of knowing are excluded from the global technological imagination. Moorosi (2024) offers a telling example: even the most advanced language models, such as GPT-4, achieve only marginal accuracy when processing African languages, reinforcing Africa's peripheral status in the data economy.

Yet, paradoxically, AI is also being heralded as a driver of African development. The African Union's 2024 Continental AI Strategy situates AI at the centre of its Agenda 2063, promoting its potential to stimulate innovation, enhance productivity, and solve pressing socio-economic challenges. Several national governments – including Kenya, Nigeria, South Africa, and Ghana – have begun crafting AI policies, investing in local research hubs, and seeking to align AI deployment with the Sustainable Development Goals. There is, therefore, a dual imperative: while AI presents opportunities for progress, it simultaneously demands critical ethical oversight to ensure that it does not exacerbate existing inequalities or replicate colonial hierarchies under the guise of technological modernity. This paper enters this contested terrain by interrogating the ethical underpinnings of AI from the perspective of African philosophy. It seeks to challenge the dominance of Western epistemologies in shaping global AI ethics, particularly the liberal individualism that characterises most mainstream frameworks. These approaches – exemplified by the OECD AI Principles, the European Commission's "Trustworthy AI" guidelines, and even UNESCO's more inclusive but still Western-leaning Recommendation on the Ethics of AI – tend to prioritise abstract rights, procedural accountability, and individual consent. While these are undoubtedly important, they often fail to account for communal notions of identity, responsibility, and justice that are central to African worldviews.

To this end, the paper advances a critical yet constructive account of algorithmic ethics rooted in Ubuntu philosophy. Ubuntu – a relational ethic encapsulated by the maxim "umuntu ngumuntu ngabantu" ("a person is a person through other persons") – offers a profound alternative to individualistic moral reasoning. It emphasises communal well-being, mutual respect, reciprocity, and shared responsibility, presenting a holistic framework for understanding moral agency in a technologically mediated society. By foregrounding Ubuntu, this paper does not merely aim to "localise" AI ethics for Africa but to offer a globally relevant critique of dominant paradigms. Ubuntu's emphasis on relational personhood and interconnected dignity holds transformative potential for reimagining how moral responsibility should be distributed across the lifecycle of algorithmic systems – from design and deployment to regulation and redress. The structure of the paper proceeds as follows: it begins with a review of current literature on AI ethics, highlighting the structural biases and epistemic exclusions inherent in mainstream models. This is followed by a detailed exposition of Ubuntu ethics and its implications for understanding moral responsibility in the digital age. The methodology section clarifies the conceptual and normative foundations of the paper, including its adoption of a decolonial, hermeneutic approach. The analysis and discussion section applies these insights to existing AI governance regimes and African policy initiatives, using illustrative case studies to demonstrate the feasibility of Ubuntu-inspired interventions. The conclusion offers recommendations for embedding African moral values into practical AI

governance frameworks. Throughout, the aim is to provide a rigorous and culturally attuned ethical framework that addresses the growing influence of AI in Africa and beyond.

Artificial Intelligence and Algorithmic Ethics

Artificial Intelligence (AI), as a technological paradigm, represents more than the mechanisation of human cognition; it embodies the values, priorities, and assumptions of its creators. AI systems, particularly those employing machine learning, operate by extracting patterns from vast datasets to make probabilistic predictions or decisions. These technologies underpin an increasingly wide range of activities – from facial recognition and credit scoring to predictive policing and health diagnostics – often operating in ways that are opaque, complex, and difficult to audit. What makes these systems ethically significant is not merely their capacity to automate decisions, but their tendency to reproduce and exacerbate existing power dynamics and social hierarchies. As such, AI is not only a technical challenge but also a profoundly moral and political one. The emergence of algorithmic ethics as a field reflects growing awareness of the normative implications of AI. Algorithmic ethics interrogates how automated systems affect concepts such as justice, accountability, privacy, fairness, and trust. Foundational questions include: Who is responsible when an algorithm harms someone? How can systems be made transparent and explainable? Are these systems fair, and by what measure of fairness? These inquiries are far from abstract. Numerous empirical studies have shown that AI systems often entrench discriminatory practices due to biased training data, flawed design assumptions, and socio-technical blind spots. For example, Buolamwini and Gebru (2018) demonstrated that commercial facial recognition systems exhibit much higher error rates for dark-skinned women than for light-skinned men, a disparity rooted in racially skewed training datasets and design choices made without diversity in mind.

Moreover, the ethical challenges of AI are not evenly distributed. Marginalised populations – whether by race, gender, class, or geography – are more likely to suffer algorithmic harms while enjoying fewer of the associated benefits. In the African context, this asymmetry is magnified by global technological dependency and historical legacies of colonialism. Much of the AI infrastructure in Africa is imported from the Global North, often without adequate consideration of local needs or socio-cultural contexts. For example, algorithmic systems used in African public services – such as biometric voter registration or educational placement algorithms – are frequently designed by foreign contractors, with limited transparency and little public oversight. These systems operate within governance frameworks that may lack the institutional robustness required to enforce ethical accountability. International responses to these concerns have taken the form of ethical guidelines and principles intended to steer AI development in socially beneficial directions. The OECD's Principles on Artificial Intelligence, the European Commission's Ethics Guidelines for Trustworthy AI, and the UNESCO Recommendation on the Ethics of Artificial Intelligence all aim to codify high-level norms such as transparency, non-discrimination, and human oversight. While commendable in their intent, these frameworks are largely the product of Euro-American liberal thought. They privilege individual rights and procedural accountability, often underemphasising collective values and contextual moral reasoning. This normative architecture aligns with the philosophical legacy of Enlightenment rationalism, which foregrounds autonomous individuals as the primary units of moral and legal concern. Consequently, these ethical frameworks often struggle to resonate with communitarian societies where identity and morality are relationally defined.

This disconnect is especially evident in the African context. Mainstream AI ethics frameworks tend to ignore or undervalue communal notions of autonomy, collective responsibility, and social harmony that are central to African moral philosophies. For instance, the concept of informed consent – a staple of Western bioethics and AI policy – presumes a well-defined, self-determining individual making decisions in isolation. However, in many African cultures, decisions with ethical significance are often made through communal deliberation or with deference to familial and social obligations. This incongruity signals a deeper issue: the epistemic foundations of algorithmic ethics are not universally shared, and without conscious

effort to diversify them, AI governance risks entrenching new forms of epistemic injustice. Epistemic injustice, as theorised by Fricker (2007), occurs when individuals or groups are unfairly excluded from the production, validation, or transmission of knowledge. In the AI context, this manifests in two key ways: testimonial injustice, where the insights or experiences of marginalised communities are disregarded or discredited; and hermeneutical injustice, where these communities lack the interpretive resources to make sense of their own marginalisation. For example, if a predictive policing algorithm disproportionately targets certain neighbourhoods without offering residents the tools or language to understand and contest these outcomes, then hermeneutical injustice has occurred. Likewise, when African voices are systematically absent from global AI discourse – as is often the case in influential policy circles and academic publications – testimonial injustice is perpetuated.

The practical consequences of such exclusions are significant. Algorithmic systems that fail to reflect the lived experiences and moral values of African communities are less likely to gain public trust or achieve just outcomes. Moreover, the absence of African philosophical insights impoverishes the global discourse on AI ethics, depriving it of alternative conceptions of personhood, morality, and justice that could enrich collective understandings of how to govern emerging technologies. In this light, African contributions to AI ethics should not be seen as local adaptations of universal norms, but as substantive interventions that challenge, complement, and potentially transform the global ethical landscape. In sum, algorithmic ethics must evolve beyond its Western-centric origins to become genuinely inclusive and globally responsive. This entails recognising the situatedness of moral values and the importance of cultural context in shaping what counts as ethical. For African societies, this means reclaiming moral agency in the digital sphere by articulating ethical frameworks that resonate with indigenous philosophies and social realities. Ubuntu, as will be argued in the following section, offers a compelling foundation for such a reorientation. It provides not only a set of moral principles but a radically different ontology of the human being – one rooted in relationality, interdependence, and shared responsibility – which can guide the ethical governance of AI in ways that are both locally grounded and globally significant.

Ubuntu Ethics and Moral Responsibility

Ubuntu, a moral-philosophical framework indigenous to many Southern African cultures, offers a rich and coherent alternative to the dominant ethical paradigms that have shaped global AI governance. The term itself, derived from Nguni Bantu languages such as Zulu and Xhosa, is often rendered as "umuntu ngumuntu ngabantu," meaning "a person is a person through other persons." At its core, Ubuntu affirms that human existence and moral agency are fundamentally relational. One becomes a person not in isolation, but through ethical engagement with others. This relational ontology stands in stark contrast to the individualistic anthropology embedded in much of Western moral theory, which conceives of the individual as a self-sufficient rational agent prior to and independent of social bonds. Ubuntu's moral universe is structured by values such as communal solidarity, mutual care, harmony, reciprocity, and respect for human dignity. These are not merely aspirational ideals but actionable principles embedded in African customary law, conflict resolution, health care, education, and social rituals. Within an Ubuntu framework, moral responsibility is never solely about individual accountability. It is always distributed, shared, and reflexive – each person is responsible for nurturing the well-being of others, just as the community is responsible for upholding the integrity of the individual. This reciprocity of care is central to African ethical traditions and offers profound implications for the governance of emerging technologies.

When applied to AI, Ubuntu ethics invites a fundamental rethinking of what counts as responsible design, deployment, and regulation. Under Ubuntu, the ethical evaluation of AI systems does not focus solely on whether an individual user's rights have been violated, but rather on how such systems affect the fabric of community life. For example, an algorithm used for resource allocation in healthcare would not only be judged by its statistical accuracy but by its impact on trust, cohesion, and equity across social groups. This communitarian criterion contrasts sharply with utilitarian metrics or rights-based indices that dominate Western

approaches. The application of Ubuntu to technological ethics is not merely speculative. Emerging research shows how African epistemologies can be operationalised in AI contexts. Odero, Nderitu, and Samuel (2024) offer a practical model wherein AI diagnostic tools are co-designed with community input, including village elders, patients, and local health workers. This approach ensures that algorithms do not simply impose technocratic efficiency but resonate with communal values and priorities. Grancia (2024) similarly demonstrates how an Ubuntu-informed approach in healthcare AI promotes relational accountability – a shared responsibility among developers, users, and affected communities to monitor and mitigate bias. Such models foreground communal consent, interpretive inclusion, and participatory governance, offering procedural mechanisms that contrast markedly with the Western emphasis on individual data subject rights and legalistic redress.

Moreover, Ubuntu reframes the very concept of autonomy. In Western liberal traditions, autonomy is typically equated with the ability to act independently and make decisions free from external interference. In Ubuntu, autonomy is relational: it is the capacity to flourish within and through one's relationships. Ewuoso (2021) elaborates this idea in the context of bioethics, arguing that relational autonomy acknowledges the interdependence of individuals and the role of social context in shaping agency. This is crucial for AI ethics, where autonomy is often framed in terms of user control over data or algorithmic decision-making. An Ubuntu-based model would instead prioritise communal deliberation and culturally informed standards of agency, reframing privacy, consent, and choice as socially embedded phenomena. Additionally, Ubuntu offers insights into modes of justice. While Western legal systems tend to be adversarial and retributive, Ubuntu-based systems are restorative and reconciliatory. Justice, under Ubuntu, aims not simply to determine fault and administer punishment, but to restore harmony among affected parties and reintegrate wrongdoers into the community. Van Norren (2023) argues that such a harmonising ethos could inform how we respond to algorithmic harm. Rather than relying solely on courts or formal grievance mechanisms, communities could establish forums for restorative dialogue, involving both developers and those impacted by AI decisions. These participatory spaces could facilitate not only accountability but ethical learning and communal healing, thus embedding moral responsibility within a broader social ecology.

Comparative philosophical insights further illuminate Ubuntu's distinctiveness. While there are conceptual overlaps with other non-Western traditions – for example, Confucian ethics in East Asia or Andean relational cosmologies in Latin America – Ubuntu stands out in its unique fusion of metaphysics and morality. It posits not just that people ought to care for one another, but that their very being is constituted by this care. This metaphysical depth enables Ubuntu to resist the reductive tendencies of algorithmic thinking, which often abstracts individuals into data points and decisions into computations. By asserting that moral worth arises from relational embeddedness, Ubuntu challenges the dehumanising logics of techno-solutionism and offers a normative compass for technology grounded in human dignity. However, Ubuntu's promise is not without limitations or risks. One critique is that communal norms can sometimes suppress individual dissent or reinforce existing hierarchies, especially where age, gender, or social status determine voice and influence. As such, a reflexive Ubuntu ethics must be open to internal critique and informed by feminist and egalitarian considerations. In addition, questions arise about the scalability of Ubuntu in pluralistic, urbanised, and digitally networked societies where traditional communal structures are less pronounced. Yet, as several scholars have argued, Ubuntu is not static or nostalgic; it is a dynamic moral grammar capable of evolving to meet new challenges while preserving its relational essence.

In conclusion, Ubuntu offers a compelling ethical vision that challenges the atomistic assumptions of mainstream algorithmic ethics. It calls for a shift from individual-centred to community-centred thinking, from abstract proceduralism to context-sensitive judgment, and from technocratic control to shared moral responsibility. In doing so, it provides not only a critique of Western models but a substantive alternative – one that can enrich global conversations on AI ethics and guide locally grounded practices. As the next section will

explore, this ethical reorientation is already underway in academic discourse, policy initiatives, and applied research across the African continent.

Literature Review

The literature on artificial intelligence and ethics has grown rapidly in recent years, reflecting widespread concern about the social and moral implications of algorithmic systems. Scholars from various disciplinary backgrounds – including philosophy, law, computer science, and sociology – have contributed to an evolving corpus that seeks to understand how AI technologies affect justice, equity, and human rights. While this body of work has produced valuable insights, it remains skewed in terms of geographic representation and philosophical orientation. The vast majority of published research continues to emerge from Euro-American contexts, with relatively little integration of African perspectives or indigenous moral traditions. This epistemic imbalance perpetuates the global asymmetries in knowledge production and contributes to what has been termed "epistemic coloniality" in AI ethics. Foundational studies by scholars such as Buolamwini and Gebru (2018) have exposed the racial and gender biases embedded in commercial AI systems, particularly facial recognition technologies. Their empirical work demonstrated that systems trained predominantly on lighter-skinned male faces exhibit far higher error rates when classifying the faces of darker-skinned women, highlighting how data bias translates into discriminatory outcomes. Noble (2018), in *Algorithms of Oppression*, similarly critiques how search engines perpetuate racial and gender stereotypes, arguing that seemingly neutral algorithms reinforce systemic inequalities. These studies provide a crucial starting point for ethical inquiry, but they also foreground the need to interrogate the socio-political structures and historical forces that shape AI systems.

Amidst these critiques, there is a growing recognition that purely technical fixes – such as diversifying training datasets or implementing fairness constraints – are insufficient. As Birhane (2021) and Mohamed, Png, and Isaac (2020) argue, such measures often fail to address the deeper structural and epistemic roots of algorithmic harm. Instead, these authors advocate for relational and decolonial approaches to ethics that situate AI within broader systems of power and historical context. This perspective is especially salient for Africa, where the deployment of AI often reproduces the patterns of extractivism and dependency that characterised colonialism and its aftermath. African scholars have increasingly responded to this challenge by articulating alternative ethical frameworks grounded in indigenous philosophies. Gwagwa, Orr, and Nkala (2022) argue that Ubuntu offers a viable basis for rethinking AI ethics in ways that prioritise communal dignity, shared responsibility, and social harmony. In a similar vein, Mhlambi and Tiribelli (2023) develop a decolonial AI governance model rooted in Ubuntu, contending that relational autonomy should replace liberal individualism as the organising principle of ethical evaluation. These interventions mark a critical shift from reactive critique to proactive theorisation, offering a normative foundation that is both locally resonant and globally relevant.

Other scholars focus on the operationalisation of African ethics in data governance and AI design. Barrett et al. (2025) propose a “seven-principle” African data ethics framework, encompassing values such as dignity, justice, solidarity, reciprocity, accountability, sustainability, and subsidiarity. These principles offer practical guidance for policymakers and developers, encouraging the translation of ethical ideals into concrete institutional practices. Mengesha, Belay, and Adams (2024) advance this agenda with their Discovery-Translation-Verification (DTV) framework, which outlines a participatory methodology for contextualising AI ethics in African settings. The DTV model stresses the need to engage community stakeholders, translate moral principles into technical specifications, and verify alignment through ongoing dialogue and feedback mechanisms. A parallel body of work engages with the concept of epistemic injustice in AI systems. Fricker’s (2007) influential theory distinguishes between testimonial injustice – when someone is discredited as a knower due to identity prejudice – and hermeneutical injustice – when structural gaps in collective interpretive resources hinder marginalised groups from making sense of their experiences. Scholars such as

Stewart, Cichocki, and McLeod (2022) apply these ideas to social media algorithms, showing how targeted advertising and content curation exacerbate testimonial exclusion. Milano and Prunkl (2024) demonstrate how algorithmic profiling contributes to hermeneutical injustice by obscuring the criteria by which individuals are classified and judged. These critiques reveal that AI does not simply process information but shapes what counts as knowledge and whose experiences are rendered intelligible.

Within African contexts, the implications of epistemic injustice are profound. AI systems developed in the Global North often fail to account for African languages, social institutions, and epistemologies. Moorosi (2024) reports that large language models such as GPT-4 perform poorly in African linguistic contexts, with accuracy rates as low as 10–20% for major languages like Hausa. This exclusion from the linguistic and cultural substrate of AI systems constitutes a form of digital erasure, reinforcing Africa's peripheral status in global data economies. Ofosu-Asare (2024) describes this as a form of "cognitive imperialism," whereby foreign systems impose epistemic frameworks that displace indigenous ways of knowing. These critiques underscore the need for ethical models that not only mitigate harm but also affirm African intellectual sovereignty. At the intersection of theory and policy, African nations are beginning to formulate responses. The African Union's 2024 Continental AI Strategy emphasises the importance of developing context-sensitive ethical frameworks and promoting digital sovereignty. Countries like South Africa and Kenya have incorporated Ubuntu principles into national AI discussions, while Nigeria's National Information Technology Development Agency (NITDA) has initiated guidelines for responsible AI use. However, as Stewart et al. (2022) note, these efforts remain fragmented and under-resourced, lacking the institutional capacity and regulatory authority to enforce ethical standards at scale. Consequently, much of the work of advancing African AI ethics falls to academic and civil society actors, who continue to push for inclusive, participatory, and decolonial approaches to technological governance.

In sum, the literature reveals an emerging but underdeveloped field of African algorithmic ethics. While the dominance of Euro-American frameworks remains pronounced, a growing body of African scholarship is challenging this hegemony by articulating alternative moral visions grounded in Ubuntu and other indigenous philosophies. These contributions are not merely regional correctives but substantive engagements with global debates on fairness, justice, and responsibility in AI. They invite a reconfiguration of algorithmic ethics – one that is attentive to history, responsive to culture, and committed to epistemic plurality.

Methodology

This study adopts a qualitative, conceptual methodology rooted in normative analysis and hermeneutic reflection. Rather than empirically testing hypotheses or quantitatively modelling AI systems, the paper aims to interrogate the moral assumptions and epistemological foundations of algorithmic ethics from the vantage point of African philosophy. This orientation is not only suitable but essential when the object of inquiry is the value-ladenness of technology, the marginalisation of indigenous worldviews, and the ethical plausibility of alternative frameworks such as Ubuntu. Accordingly, the research design privileges theoretical depth, cultural specificity, and moral clarity over empirical generalisation. The central methodological approach employed is a normative-constructivist one: the study explores what ought to be the case in AI governance, based on a critical evaluation of existing frameworks and the ethical affordances of African philosophical traditions. Normative analysis entails not only assessing current ethical models for their internal coherence and applicability but also proposing alternative principles that better align with justice, inclusivity, and epistemic integrity. This task involves interpreting Ubuntu as a living moral philosophy capable of guiding technological governance in a way that is culturally consonant and globally constructive. The ethical reasoning here proceeds from African communitarian premises, in contrast to liberal individualism, and aims to articulate both the content and implications of a relational moral ontology.

In order to translate these insights into a rigorous and context-sensitive critique, the study employs hermeneutic analysis, which foregrounds interpretation, meaning-making, and the role of cultural narratives in shaping moral understanding. A hermeneutic methodology is especially apt when dealing with African philosophy, which is largely oral, narrative-driven, and embedded in communal life. As Wiredu (1996) and Gyekye (1997) have noted, African moral thought is best understood not as a system of abstract rules but as a holistic practice of lived relationality, whose meaning emerges through participation, reflection, and shared experience. Thus, this study reads Ubuntu not merely as a set of aphorisms but as a coherent ethical framework whose norms and values must be interpreted in relation to their socio-cultural and historical contexts. Complementing the hermeneutic framework is a decolonial ethos, which insists on disrupting epistemic hierarchies and validating local knowledge systems. This orientation follows scholars such as Ofosu-Asare (2024), who argue that AI ethics must be decolonised by embedding indigenous epistemologies into the structures of knowledge production and policy design. A decolonial methodology resists the universalisation of Eurocentric norms and seeks instead to pluralise the sources of ethical reflection. In practice, this means that the study engages primarily with African sources, both classical and contemporary, and interprets global AI ethics guidelines through an African lens. It also pays attention to the socio-political structures – including histories of colonialism, marginalisation, and resistance – that shape Africa's engagement with digital technologies.

To provide practical grounding, the paper adopts the Discovery-Translation-Verification (DTV) framework developed by Mengesha, Belay, and Adams (2024). This three-stage process serves as both a conceptual heuristic and a procedural guide for contextualising AI ethics. In the *discovery* phase, the study identifies core ethical principles within Ubuntu, such as relational autonomy, communal responsibility, and social harmony. The *translation* phase involves interpreting these principles in ways that are intelligible to AI designers, policymakers, and regulators. This includes specifying how values like "communal consent" might be integrated into data governance, or how "restorative justice" could shape redress mechanisms for algorithmic harm. Finally, the *verification* phase entails evaluating these translated principles against the standards of coherence, cultural legitimacy, and practical applicability, using illustrative case studies and theoretical reflection. Reflexivity also plays an important role in the research process. The positionality of the author – as a scholar embedded within African philosophical and academic traditions – informs the selection of sources, the interpretation of ethical concepts, and the normative judgments advanced. This standpoint is not a limitation but a methodological asset, enabling the work to foreground insider perspectives and challenge dominant narratives from within an engaged, situated epistemology. Moreover, the use of cross-disciplinary literature – spanning philosophy, AI ethics, African studies, and policy analysis – ensures that the analysis remains intellectually robust and accessible to multiple audiences.

In summary, the methodology of this study integrates normative ethics, hermeneutic interpretation, and decolonial critique to produce a philosophically grounded and culturally situated contribution to AI ethics. By treating Ubuntu as both a moral framework and an epistemic intervention, the research offers a rigorous and context-sensitive alternative to prevailing ethical paradigms. It sets the stage for the subsequent analytical section, which applies this framework to evaluate existing AI governance practices and proposes pathways for institutionalising Ubuntu in digital policy across Africa.

Analysis and Discussion

The critical analysis of global algorithmic ethics frameworks reveals a persistent marginalisation of African epistemologies and values. Dominant models – such as those articulated by the OECD, UNESCO, and the European Commission – are primarily underpinned by liberal democratic values and philosophical traditions rooted in the European Enlightenment. These frameworks prioritise individual autonomy, procedural fairness, and transparency, framing ethical AI in terms of rights-based protections, legal accountability, and risk mitigation. While these values have universal appeal in theory, they are under-theorised

with respect to cultural pluralism and fail to accommodate worldviews where moral agency is communally situated. The result is an ethical architecture that, though nominally global, remains epistemically parochial. Ubuntu offers a necessary corrective to this asymmetry. Its emphasis on relational autonomy, shared moral responsibility, and communal flourishing foregrounds dimensions of justice that individualistic models often obscure. As Van Norren (2023) argues, Ubuntu is not merely a cultural idiosyncrasy but a robust moral theory with translatable implications for technology governance. For example, where Western ethics might assess an AI system based on its respect for personal privacy, an Ubuntu-informed perspective would also inquire into the system's impact on collective trust, social cohesion, and the dignity of the community. This shift in focus realigns AI ethics away from atomistic safeguards and toward relational well-being.

One illustrative domain is data governance. In many African societies, the notion of "consent" – a cornerstone of Western data ethics – is not primarily individual but communal. For instance, in traditional African medicine and land tenure systems, decisions are often made collectively by elders or councils, reflecting a normative belief that individual choices are embedded within broader social responsibilities. Applying this to AI, a system that collects biometric data for public health purposes must seek not only individual permission but community endorsement, ensuring alignment with shared cultural and ethical norms. The African Union's 2024 AI Strategy hints at this possibility by advocating for "context-sensitive" consent models, although operational frameworks remain inchoate. This communitarian orientation also alters our understanding of accountability. In Western AI governance, responsibility is typically ascribed to identifiable actors – designers, developers, regulators – in a manner that presupposes linear causality and discrete agency. Ubuntu, by contrast, views moral responsibility as distributed across relational networks. When an algorithm causes harm – for instance, by denying someone a loan or misidentifying a face – the question is not simply who authored the code, but how the socio-technical system as a whole failed to honour the communal duty of care. This perspective encourages forms of collective accountability, such as stakeholder assemblies or community oversight boards, which can deliberate on harms in ways that are restorative rather than adversarial.

Some countries have already begun to implement AI governance models with Ubuntu elements. In Kenya, the National Commission on AI has endorsed participatory mechanisms in policy consultations, engaging civil society actors, religious groups, and traditional leaders in discussions about data privacy and AI use in health care. This inclusive strategy mirrors Ubuntu's dialogical ethos, wherein moral legitimacy is derived not from top-down imposition but from consensus and mutual recognition. Likewise, South Africa's Presidential Commission on the Fourth Industrial Revolution has cited Ubuntu explicitly in its recommendations, noting the importance of embedding human dignity and social harmony into digital policy design. However, these initiatives are not without their shortcomings. As Stewart, Cichocki, and McLeod (2022) observe, many policy frameworks that invoke Ubuntu do so symbolically rather than substantively. References to African values are often appended as rhetorical flourishes, without mechanisms to translate them into technical standards, legislative mandates, or enforcement procedures. There is thus a risk of ethical tokenism – the superficial inclusion of local concepts that do not meaningfully shape institutional behaviour. To guard against this, African states and regional bodies must develop formal instruments that operationalise Ubuntu principles in measurable and enforceable ways.

Concrete proposals have begun to emerge from African scholars and technologists. Barrett et al.'s (2025) seven-principle African data ethics framework, for example, offers a viable template for codifying Ubuntu values into actionable governance criteria. The framework outlines operational benchmarks for dignity, justice, solidarity, and accountability in data collection, processing, and application. These benchmarks can be adapted into legal codes, audit protocols, and public engagement guidelines. Similarly, the DTV model (Mengesha et al., 2024) emphasises iterative consultation with community stakeholders at every stage of algorithmic development – from initial design to post-deployment evaluation – thereby institutionalising ethical reflexivity and cultural sensitivity. The health sector offers a

compelling case study. In 2023–24, a consortium of African medical researchers and software engineers piloted an AI diagnostic tool for rural clinics in Ghana and Uganda. Rather than imposing a top-down solution, the team adopted an Ubuntu-inspired model of development: local nurses, patients, and spiritual leaders were engaged in design workshops, and ethical decisions about data use were made through consensus-building sessions. The resulting tool not only improved diagnostic accuracy but also enhanced trust in the technology, as communities saw their values reflected in the system's logic and structure. This participatory model demonstrates that Ubuntu is not merely a philosophical abstraction but a practical guide to equitable technological design.

Another significant site of Ubuntu's relevance is algorithmic redress. Western legal traditions generally favour compensatory or punitive remedies, such as monetary damages or regulatory fines. While effective in certain contexts, these approaches may not facilitate the communal healing necessary after algorithmic harms. Ubuntu suggests a different model: restorative justice. For example, when a predictive policing algorithm disproportionately targets a marginalised group, the appropriate response may involve structured dialogues between developers, affected citizens, local leaders, and civil society groups. These dialogues would aim not just at restitution but at restoring moral equilibrium and reconstituting trust. Such processes align with traditional African dispute resolution mechanisms and offer a culturally congruent approach to algorithmic justice. Yet challenges remain. Foremost among them is the entrenched dominance of global North institutions in setting the agenda for AI development and ethics. International regulatory bodies, technical standard-setting organisations, and research consortia continue to privilege Western norms, languages, and epistemologies. Even when African countries are consulted, their input is often marginalised or co-opted. This structural exclusion constitutes what Ofosu-Asare (2024) terms "cognitive imperialism" – the monopolisation of technological meaning-making by epistemic centres in the North. Without deliberate efforts to decentralise and democratise AI governance, Ubuntu's transformative potential will remain unrealised.

A further issue concerns the scalability of Ubuntu ethics in highly urbanised, digitally interconnected environments. While Ubuntu emerged within small-scale, face-to-face communities, many African societies today are marked by fragmentation, inequality, and globalised cultural flows. This raises important questions: Can Ubuntu adapt to the complexities of digital life without losing its ethical coherence? Can it inform algorithmic accountability in a world where decisions are made by inscrutable neural networks and transnational corporations? These challenges demand careful theoretical work and institutional experimentation. Nonetheless, as Ewuoso (2021) contends, Ubuntu is a dynamic tradition, not a static relic. It can evolve, extend, and reframe its moral insights in ways that address contemporary realities without abandoning its core commitment to relational dignity. In summary, Ubuntu provides a comprehensive moral grammar for rethinking algorithmic ethics in Africa and beyond. It challenges prevailing models of responsibility, autonomy, and justice, replacing them with a relational ethics that is both contextually grounded and normatively rich. While some African states and scholars have begun to explore its potential, much work remains in translating Ubuntu from philosophical aspiration to regulatory and technical practice. The next and final section outlines specific recommendations for doing so and proposes institutional pathways for embedding African values in the global ethics of artificial intelligence.

Conclusion and Recommendations

The analysis throughout this paper has demonstrated that artificial intelligence, while offering unprecedented opportunities for development, innovation, and efficiency, also presents deep ethical and philosophical challenges—particularly for African societies. These challenges are not merely technical or regulatory; they are fundamentally epistemic and moral. The prevailing global frameworks for AI ethics, while noble in aspiration, are predominantly shaped by Western liberal traditions that valorise individual rights, procedural safeguards, and universalist norms. Such frameworks often fail to accommodate the moral grammar of African communities, where values such as relational personhood, communal responsibility, and ethical

interdependence are paramount. In this context, Ubuntu philosophy emerges not merely as an indigenous alternative but as a globally significant ethical paradigm capable of reconfiguring the foundational principles of AI governance. Ubuntu, as a relational ethic, foregrounds the interconnectedness of human beings, the primacy of community, and the necessity of shared moral responsibility. It challenges the atomism of Western ethics by positing that human identity and agency are constituted through social relations, not in abstraction from them. This perspective has profound implications for how we conceive moral responsibility in the design, deployment, and regulation of AI systems. Rather than asking merely whether an AI respects individual rights, an Ubuntu-informed ethics asks whether it nurtures community well-being, restores broken relationships, and upholds the dignity of all stakeholders. This reconceptualisation enables African societies—and the global community—to confront algorithmic injustice not only with critique but with an affirmative vision of ethical technology. However, for Ubuntu to make a substantive impact on AI governance, it must move from philosophical discourse to institutional implementation. Symbolic invocations of African values are insufficient if they are not accompanied by concrete legal frameworks, policy reforms, technical standards, and educational initiatives. The goal is not simply to "Africanise" AI ethics as a cultural project, but to co-create a genuinely pluralistic and just technological future where African voices and moral systems help shape global norms. To this end, several interrelated recommendations are proposed.

1. Institutionalise African Ethical Frameworks in AI Governance

Governments across Africa should prioritise the development of national and regional ethical guidelines for AI that are explicitly grounded in African philosophical traditions. These frameworks should not mimic Western templates with superficial modifications but should operationalise values such as Ubuntu, solidarity, communal dignity, and restorative justice. The African Union, in collaboration with continental bodies such as the African Academy of Sciences and the Council for Scientific and Industrial Research, could play a coordinating role in drafting a continental charter on AI ethics that reflects African intellectual sovereignty.

2. Create a Pan-African AI Ethics Observatory

To ensure accountability, monitor implementation, and foster innovation in culturally grounded AI governance, a Pan-African AI Ethics Observatory should be established. This independent body could serve as a clearinghouse for best practices, conduct audits of AI systems, support policy harmonisation across member states, and amplify African research in international forums. Importantly, it should include ethicists, technologists, civil society actors, traditional leaders, and affected communities to ensure inclusive and pluralistic governance.

3. Embed Ubuntu Principles in Technical Design and Evaluation

AI developers working within African contexts should be encouraged—and where appropriate, mandated—to integrate Ubuntu principles into their design processes. This includes participatory design methodologies that involve community stakeholders in decision-making; data governance models that reflect communal consent; and redress mechanisms that emphasise reconciliation and relational accountability. Existing frameworks such as the DTV model and the seven-principle African data ethics rubric provide practical tools for embedding Ubuntu into technical workflows.

4. Strengthen Ethical Education in STEM Curricula

African universities and technical institutes should revise their curricula in computer science, engineering, and data science to include mandatory courses on African ethical thought, AI ethics, and decolonial philosophy. Students should be exposed to case studies that illustrate the consequences of algorithmic harm and be trained to evaluate their work not only for technical efficiency but for moral impact. Such education would cultivate a generation of developers and

policymakers equipped to think critically and ethically about technology in ways that are locally grounded and globally competent.

5. Facilitate Community-Led AI Research and Innovation

Donor agencies, academic institutions, and African governments should increase investment in community-led AI research that centres indigenous knowledge systems and local priorities. Funding mechanisms should incentivise research that explores how Ubuntu and other African philosophies can inform algorithmic design, data governance, and impact assessment. Collaborations between universities, traditional institutions, and grassroots organisations should be supported to ensure that research outputs are not only theoretically sound but socially embedded.

6. Promote African Leadership in Global AI Ethics Debates

African scholars, ethicists, and policymakers must be strategically positioned in global forums that shape AI governance, such as the United Nations, the OECD, IEEE, and ISO. Their participation must go beyond symbolic representation; they should serve as agenda-setters and normative architects. This requires advocacy, coalition-building, and diplomatic engagement to ensure that African perspectives are not marginalised or diluted in the pursuit of supposed universality.

7. Foster Restorative Approaches to Algorithmic Harm

Where algorithmic systems have caused harm—through misclassification, exclusion, or discriminatory outcomes—mechanisms for redress should reflect Ubuntu's commitment to restoration and harmony. Community dialogues, public apologies, and participatory restitution schemes can complement legal remedies and cultivate public trust in AI systems. These approaches resonate with African traditions of conflict resolution and offer a culturally congruent model of justice.

In conclusion, Ubuntu provides more than a regional ethical lens—it offers a transformative moral vision for AI governance in an age marked by digital interdependence and epistemic pluralism. It encourages us to reimagine responsibility not as a linear chain of culpability but as a shared commitment to human flourishing. It urges a move from technocratic fixes to moral imagination, from data extraction to epistemic justice, and from isolated rights to collective dignity. For Africa to contribute meaningfully to the global future of AI, it must assert not only its technical agency but its ethical authority. The road ahead requires courage, collaboration, and the conviction that technology, like humanity, is at its best when it honours our relational essence.

References

Barrett, L., Mahamedou, T., Onuoha, V., & Ekezie, J. (2025). *African data ethics: A seven-principle framework for responsible data governance*. *Journal of African AI Ethics*, 12(3), 45–67.

Birhane, A. (2020). Algorithmic colonisation of Africa. *SCRIPTed: A Journal of Law, Technology & Society*, 17(2), 389–408.

Birhane, A. (2021). Algorithmic injustice: A relational ethics perspective. *Patterns*, 2(2), 100205. https://doi.org/10.1016/j.patter.2021.100205

Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 1–15.

Ewuoso, C. (2021). Relational autonomy and Ubuntu in bioethics. *African Journal of Ethics*, 9(2), 33–49.

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Vayena, E. (2018). AI4People: An ethical framework for a good AI society. *Minds and Machines*, 28(4), 689–707. https://doi.org/10.1007/s11023-018-9482-5

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.

Gwagwa, A., Orr, R., & Nkala, D. (2022). The role of Ubuntu in global AI ethics discourse. *AI & Society*, 37(4), 981–994. https://doi.org/10.1007/s00146-021-01238-z

Grancia, L. (2024). Data ethics in the era of healthcare artificial intelligence in Africa: An Ubuntu philosophy perspective. *AI Ethics Journal*, 2(1), 15–32.

Gyekye, K. (1997). *Tradition and modernity: Philosophical reflections on the African experience*. Oxford University Press.

Kay, J., Kasirzadeh, A., & Mohamed, S. (2024). Generative algorithmic epistemic injustice. *arXiv preprint* arXiv:2408.11441. https://arxiv.org/abs/2408.11441

Mahamedou, T., Onuoha, V., & Barrett, L. (2024). African data ethics: A discursive framework for Black decolonial data science. *arXiv preprint* arXiv:2502.16043. https://arxiv.org/abs/2502.16043

Mengesha, Y., Belay, M., & Adams, A. (2024). Discovery, translation, and verification: An African approach to AI ethics. *Journal of African Technology Ethics*, 15(1), 11–35.

Mhlambi, S., & Tiribelli, G. (2023). Ubuntu as a decolonial ethical framework for AI governance. *AI & Society*, 38(2), 854–874. https://doi.org/10.1007/s00146-022-01376-9

Milano, S., & Prunkl, C. (2024). Algorithmic profiling as a source of hermeneutical injustice. *Philosophical Studies*, 181(1), 35–57. https://doi.org/10.1007/s11098-023-01861-y

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21. https://doi.org/10.1177/2053951716679679

Mohamed, S., Png, M. T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33, 659–684. https://doi.org/10.1007/s13347-020-00405-8

Moorosi, N. (2024, December 11). Better data sets won't solve the problem — we need AI for Africa to be developed in Africa. *Nature*. https://www.nature.com/articles/d41586-024-03678-0

Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.

Odero, B., Nderitu, D., & Samuel, G. (2024). The Ubuntu way: Ensuring ethical AI integration in health research. *Wellcome Open Research*, 9, 625. https://doi.org/10.12688/wellcomeopenres.18740.1

Ofosu-Asare, Y. (2024). Cognitive imperialism in artificial intelligence: Counteracting bias with indigenous epistemologies. *AI & Society*. https://doi.org/10.1007/s00146-024-01689-z

Polson, N. G., & Scott, J. G. (2018). *AIQ: How people and machines are smarter together*. Random House.

Ruttkamp-Bloem, E. (2023). African relational ethics and AI policy frameworks. *Ethics and Information Technology*, 25(1), 13–29. https://doi.org/10.1007/s10676-022-09641-3

Stewart, H., Cichocki, E., & McLeod, C. (2022). A perfect storm for epistemic injustice: Algorithmic targeting and sorting on social media. *Feminist Philosophy Quarterly*, 8(3/4), Article 8. https://doi.org/10.5206/fpq/2022.3-4.13424

Van Norren, D. E. (2021). The ethics of artificial intelligence through the lens of Ubuntu. In *Africa Knows Conference Proceedings*. Leiden University.

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), 14–29. https://doi.org/10.1108/JICES-05-2022-0057

Wiredu, K. (1996). *Cultural universals and particulars: An African perspective*. Indiana University Press.